NEAR-CRITICAL SPANNING FORESTS AND RENORMALIZATION

STÉPHANE BENOIST, LAURE DUMAZ, AND WENDELIN WERNER

ABSTRACT. We study random two-dimensional spanning forests in the plane that can be viewed both in the discrete case and in their scaling limit as slight perturbations of an uniformly chosen spanning tree. We show how to relate this scaling limit to a stationary distribution of a natural Markov process on a state of abstract graphs with non-constant edge-weights. This simple Markov process can be viewed as a renormalization flow, so that in this two-dimensional case, one can give a rigorous meaning to the fact that there is a unique fixed point (ie. stationary distribution) in two dimensions for this renormalization flow, and when starting from any two-dimensional lattice, the renormalization flow (ie. the Markov process) converges to this fixed point (ie. converges in law to its stationary distribution).

While the results of this paper are dealing with the planar case and build on the convergence in distribution of branches of the UST to SLE_2 as well as on the convergence of the suitably renormalized length of the loop-erased random walk to the "natural parametrization" of the SLE_2 (a recent result by Lawler and Viklund), this Markov process setup is in fact not restricted to two dimensions.

1. Introduction

Phase transition and critical phenomena are now considered from the physics point of view to be a fairly settled issue, thanks to numerous important works of these last 70 years. On the mathematical side, there now exists a couple of important discrete two-dimensional models for which one can really prove that the discrete critical system converges to a continuous scaling limit (that turns out to be conformally invariant), but many fundamental problems remain unsolved. This includes the existence and the description of scaling limits for three-dimensional models, and the understanding of the universality question (for instance: How can one prove that for a given model – say the percolation model – in a given dimension, all its critical versions behave in the same way in the scaling limit, independently of the chosen lattice?).

In [25], a fairly simple renormalization formalism is described that enables to view the conjectured scaling limits of perturbations of the critical FK-percolation models (recall that these form a family of models very closely related to Ising and Potts models) as stationary distributions of a simple Markov process on a state of discrete random weighted graphs (see [25] for more details, motivation and references). In the present paper, we explain how to implement this set-up for one of the mathematically best understood models, namely the uniform spanning tree (UST). Recall that one can interpret this model as being part of the class of critical FK-percolation models (in the limit when the parameter q governing the FK model goes to 0) and that existence of the scaling limit, its description (via SLE curves) and universality (i.e. USTs on different lattices have the same scaling limit) have been proved in two dimensions [17], and that some partial results are even available in three dimensions [12].

In some sense, our approach in the present paper is the analogous study - in the case of the UST - of some features of critical percolation that had been derived by Garban, Pete and Schramm [8, 9, 10], who described the scaling limit of what one sees when one dynamically goes through the phase transition of planar percolation (when opening at random and independently the edges or sites one after the other).

Our perturbation of the UST corresponds to a simple two-step procedure: First sample a UST and then erase some of its edges uniformly at random (in a Poissonian way). One then obtains a random forest because some edges of the tree are missing. It is a *uniformly cut uniform spanning tree.*

Alternatively, we can view the erasing procedure in the reverse way, and let the different edges of the UST appear one after the other in random order (and possibly stop a little bit before all edges have appeared). If we now observe these edges appearing one by one (using the previously described rule) without knowing a priori what the UST is going to be, one discovers only progressively which edges are eventually going to be in the UST. The first important and simple observation is that this dynamics is in fact a simple Markov process on a set of subgraphs (i.e. on the set of forests) of the original graph, with jumps of this Markov process corresponding to the opening of an edge (i.e. to the merging of two trees into a single larger tree).

Let us spend a few lines to explain in more detail this simple discrete dynamical Markovian construction of the UST in a finite connected graph G with N vertices. At step one, sample a UST in G, and choose uniformly at random one of the $N-1$ edges of this UST, that we call e_1. Now, at time 1, we define G_1 to be the graph with vertices those of G, but with just one open edge, e_1. Next, we want to choose a second edge. The conditional law of the UST given that e_1 has been chosen in the first step is just the distribution of the UST conditionally on the event that it contains e_1 (this is just because for each tree that contains e_1, the probability to choose e_1 in this first step is the constant $1/(N-1)$). Hence, in order to choose the second edge e_2, we can sample a UST in the new graph obtained when merging the extremities of e_1 and then choose an edge uniformly at random in this tree. After $N-1$ such steps, one has in this way constructed a UST $\mathcal{T}$ on G with edges $e_1, \ldots, e_{N-1}$. It is then immediate to check that this construction of the UST is the time-reversal of the procedure where we first sample the UST $\mathcal{T}$, and then erase its edges one after the other uniformly at random.

This discrete-time procedure has a simple continuous-time counterpart. Since the number of edges in a spanning tree is anyway a graph-dependent constant, one can choose the rate at which an edge e between x and y opens at a given time t to be the probability that e will eventually belong to an UST in the graph G_t obtained from G by contracting all its edges that have been already open before time t. Note that this rate is just the effective resitence of the edge e in the graph G_t (ie. it is the current that flows through e when a total of one unit of current is sent from the tail of e to the tip of e, see for instance [20]).

It is easy to generalize this continuous-time Markov process construction of the UST to infinite graphs, for instance to $\mathbb{Z}^d$. Then, when time is getting very large, while locally (say near the origin) the structure of the graph converges to that of the UST (or the uniform spanning forest (USF) if $d \geq 5$), on a very large scale L, there are still a lot of edges of the final tree (or forest) that have not yet been opened, so that the clusters that one sees are all of size much smaller than L. And if one chooses an appropriate scale $L(t)$ (a procedure referred to as finite-size scaling), then the picture looks a little bit like the scaling limit of the UST that has been cut into multiple pieces in a Poissonian way. Hence, modulo appropriate rescaling, the picture when $t \to \infty$ looks like a "near-critical forest".

In order to study the behavior of this Markov process in its scaling limit, we will be to introduce a slightly different setting. Each forest F in our original graph G is naturally associated to a more abstract graph $S(F)$ with edge-weights, the structure graph of F:

- Clusters c of F corresponds in a one-to-one way to sites $s(c)$ of $S(F)$.
- When two clusters c and c' are not adjacent, there is no edge joining $s(c)$ and $s(c')$.

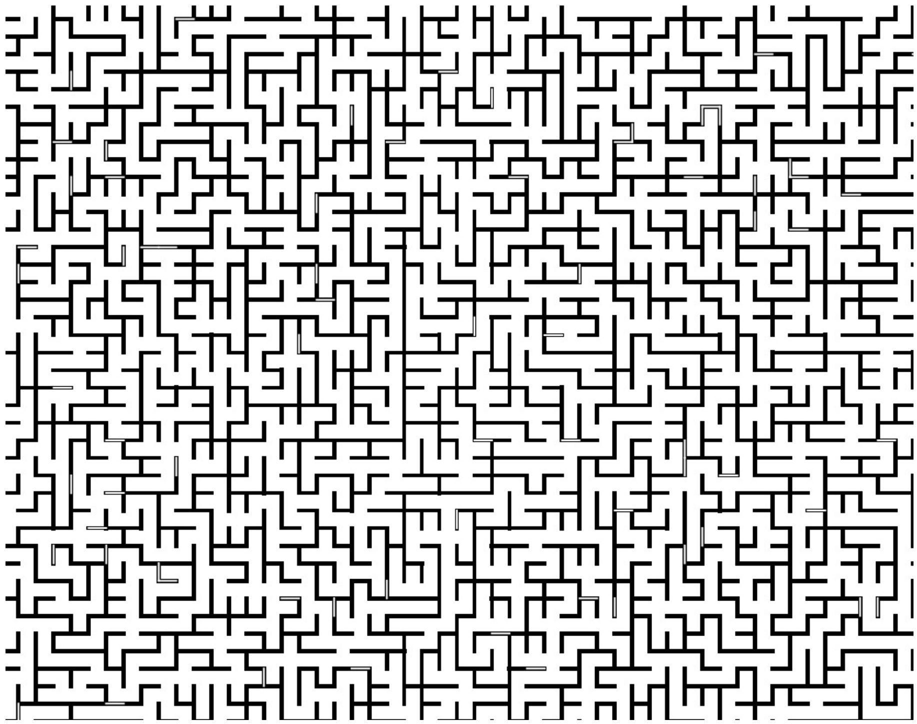

FIGURE 1. Cutting the UST in a Poissonian way (the white edges are removed from the tree)

- When two clusters c and c' are adjacent, then $s = s(c)$ and $s' = s(c')$ are joined in $S(F)$ by an edge (s, s') with weight $w(s, s')$ given by the number of edges (of the original graph G) that connect c to c'.

The previously described Markovian dynamics $(F_t, t \geq 0)$ on the state of forests induces a Markovian dynamics $(S(F_t), t \geq 0)$ on this state of weighted graphs (because the only information about F_t that is used to describe the future evolution of the forest can be encapsulated in $s(F_t)$). The time-evolution for $S(F_t)$ then corresponds to the merging of neighboring sites s' and s (i.e. collapsing of the edge (s, s') between them) that occurs at a certain rate depending on all the weights w (i.e. it depends on the entire graph $S(F_t)$). When one collapses s and s' into a site ss', then the new edge-weights $\tilde{w}$ are simply given by $\tilde{w}(ss', s'') = w(s, s'') + w(s', s'')$, while the edge-weights corresponding to edges that are not adjacent to ss' are left untouched.

Note that if F_0 has finitely many sites, then for any large enough time, $s(F_t)$ is just a graph with one single point and no edge (as it corresponds to the image of a tree under S).

We are now finally ready to describe in loose words the content of the main results (Theorems 4 and 5) of the present paper :

(1) We will first see that the definition of our Markov process on discrete structure graphs can be extended to a space of graphs with unbounded degrees, that one can view as the space of scaling limits of discrete structure graphs. The core of the paper is then to prove that the scaling limit of the Markov process (on discrete graphs) is indeed this Markovian continuous process (it is not clear whether the scaling limit of a Markov process is Markov). This will already build on the description of the scaling limit of UST via SLE, and on the predicted convergence of the renormalized length of these branches to their continuous counterparts.

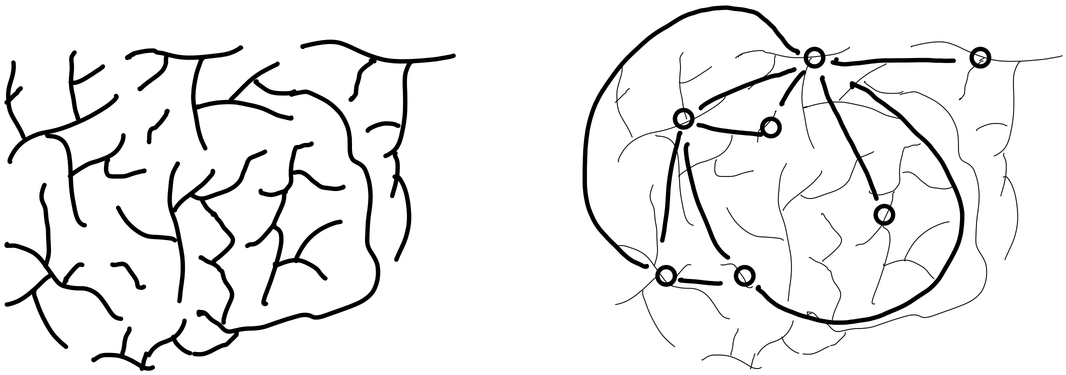

FIGURE 2. From the forest to the structure graph (sketch)

(2) Then, we will observe that in the case where one focuses on the UST in the whole plane, one gets a probability measure on this space of configurations that will be invariant under a rescaling of this Markov process. This stationary measure can be interpreted as the scaling limit of near-critical planar UST (it describes random forests). Hence in the present case, one has described a setting in which renormalization makes complete rigorous sense: when one starts from a two-dimensional lattice and runs the Markov process, the graph converges locally (in distribution) to the UST, and when one looks at the graph from further and further away, it converges to the fixed point of this Markov process, that describes the near-critical scaling limit (rather than the scaling limit itself).

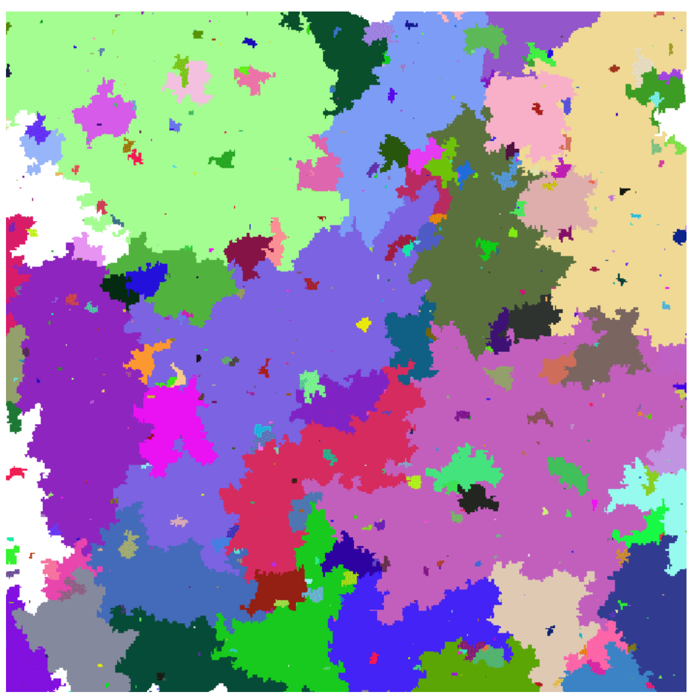

FIGURE 3. Simulation of part of the uniformly cut UST

Here is a list of some of the main technical features and tools that we shall use:

- We will use the framework introduced by Schramm [22] in order to describe the set in which our discrete objects (the UST, the near-critical forests) and their scaling limits live in: One encodes the limit of these near-critical forests to be the (countable) family of its "continuous backbone branches" (corresponding to the limit of the macroscopic branches of the cut UST).
- On this set of continuous forests, we will then define the dynamics. While the discrete dynamics are clearly Markovian, it is not obvious at all that the continuous process is (as some information may have disappeared in the scaling limit). This is the same key-problem as in the case of near-critical percolation studied in [8, 9, 10]. In order to prove this, we need a careful analysis of the discrete to continuous limiting procedure, and we shall use some stochastic comparisons between the evolutions of various graphs under our dynamics.
- We rely on the convergence of the branches of the UST (i.e. loop-erased random walks converge to SLE_2, as proved in Lawler, Schramm, Werner [17]) but one also needs to control the clocks of our dynamics i.e. the number of edges on these branches, as they control the time-evolution. For this, we will in fact use the convergence of the loop-erased random walk towards the SLE_2 in this "natural parametrization" for the uniform topology (this result will be recalled in the next section), due to Johansson-Viklund and Lawler [19] (and builds on earlier work of these authors with Benes [4]).

The paper will be structured as follows:

- In Section 2, we first recall some features of UST's, briefly define Schramm's framework for scaling limits, and investigate the scaling limit of the cutting of UST's in bounded domains.
- In Section 3, we study the time-reversal of the cutting dynamics seen on "structure graphs", and state our first main result, i.e. that this time-reversal is Markovian. We then explain why the whole-plane version of these results can be interpreted in terms of a renormalization flow fixed point.
- In Section 4, we prove the technical lemmas on discrete UST events, that are needed in the previous proofs.
- In the appendix, we use the results of [19] in order to derive the actual facts about convergence of length to natural parametrization in the settings that we need.

Let us conclude this introduction with a few words about "near-critical" models, in order to stress the fact that the near-critical forests that we construct and describe here are not exactly part of the FK-percolation family (this feature also appears in the general setup described in [25]). Recall that the word critical (in critical models) here usually refers to the fact that one considers a one-parameter family of lattice models, and that there is a phase-transition for this "critical" value of the parameter. However, there are often more than one parameter that one can play with in order to perturb the discrete model, and therefore several possible near-critical models.

On a finite-graph, it is well known that the law P^0 of the uniform spanning tree can be viewed as the limit when $p \to 0+$ and $q = o(p)$ of the random cluster (or FK)-measure $P_{p,q}$ (indeed, the fact that $q \to 0$ faster than p ensures that most of the mass of $P_{p,q}$ sits on the configurations with just one connected component, and the fact that $p \to 0$ ensures that the system uses the minimal amount of edges).

When $q \to 0$ and $p > 0$ remains fixed, the measure becomes simply $P_{p,0+}$ which is critical percolation conditioned to have exactly one connected component. On the other hand, when $q \to 0$ and p is of the same order as q, then the limit will be supported on forests (i.e. collections of trees). More precisely, when $p = q$, the $q \to 0$ limit is the uniform measure on forests and when $p = \alpha q$, the limit measure is the percolation measure of parameter $\alpha/(1+\alpha)$ conditioned on the non-existence of open circuits. This leads (via finite-site scaling, tuning $\alpha(N)$ appropriately, and letting $q \to 0$ and $N \to \infty$) to a continuous model, that corresponds to a near-critical continuous uniform spanning forest, and is a perturbation of the continuous UST. However, the object obtained via such

a construction will differ from the one that we study in the present paper (and that is obtained from first sampling the UST and then cutting some edges uniformly at random). Indeed, the discrete measures $P_{\alpha 0+,0+}$ assign the same probability to different forests that have the same number of trees, whereas in our cutting perturbation of the UST, the weight of a configuration depends in a non-trivial way on the lengths of the boundaries between the trees in the forest (as they indicate how many possible ways there were to construct the forest by cutting a tree at random).

2. UST and UST limits

2.1. **General UST Background.** Let us very quickly browse through some of the standard UST features and definitions that we will use.

The uniform spanning tree (UST) $\mathcal{T}(G)$ on a finite connected graph G is a random subgraph of G that has been uniformly chosen among those connected subgraphs that contain all vertices of G, and are cycle-free. If G is an infinite graph, one can define a similar object $\mathcal{T}(G)$, the free uniform spanning forest or USF (see e.g. [6]), as the weak limit of USTs on G_n, where G_n is any increasing exhaustion of G by finite connected subgraphs. Depending on the infinite graph, this uniform spanning forest can be almost surely a tree, or not. In $\mathbb{Z}^d$ for $d \leq 4$, the free spanning forest is actually a.s. a tree (and called a UST as well).

The notion of UST can be extended to weighted graphs: Let $G = (V, E)$ be a finite graph and $c : E \to \mathbb{R}_+$ denotes its weights. Then the weighted spanning tree is the probability measure on the set of all spanning trees such that the probability to choose a tree T is proportional to $\prod_{e \in T} c(e)$. If G is an infinite weighted graph, one can define the *weighted free spanning forest*, a probability measure on the subgraphs of G, as the weak limit of the weighted spanning tree on G_n (where G_n is any connected exhaustion of the weighted graph G). This definition in fact works even if the graph is not locally finite (i.e. sites are allowed to have infinitely many neighbors, and the sum of the incoming weights is even allowed to be infinite). Similarly, depending on G, this weighted free spanning forest can almost surely be a tree or not.

Suppose now that T is a spanning tree of the graph G and that V is a finite set of vertices G. We will denote by T_V the minimal connected subgraph of T containing V. If F is a forest (a disjoint union of trees) of G, then we define F_V as the reunion of the subtrees generated by V on all the connected components of F.

Wilson [27] provided an algorithm to sample from the UST measure on a finite graph G, by iteratively generating branches as loop-erased random walks on the graph G: Enumerate the vertices of your graph G : $x_0, x_1, \cdots, x_N$. At the n-th step of the algorithm we have constructed a tree $T_n \subset G$ that will eventually turn out to be $\mathcal{T}_{x_0,\ldots,x_n}$. Start with a single point $T_0 = x_0$. To build T_n, run a simple random walk X_n on G started from x_n and stopped upon hitting T_{n-1}. Consider the (chronological) loop-erasure γ_n of X_n, and let $T_n = T_{n-1} \cup \gamma_n$. Then, the final tree T_N has the law of a UST on G.

It is well-known that Wilson's algorithm can be extended to (locally finite) infinite graphs such as $\mathbb{Z}^d$, as well as to weighted graphs (one just needs to replace the simple random walk by a random walk with non-constant conductances). This generates a random infinite forest, known as the wired spanning forest. In $\mathbb{Z}^d$ or in graphs that are obtained from $\mathbb{Z}^d$ by contracting or erasing some edges, the free USF and the wired USF coincide, see [6].

At some points in the paper, we will use coupling results between USTs in various domains (this type of result is in fact instrumental in deriving the existence and properties of some of the objects mentioned above, such as the free USF).

Let us first recall ([6, Corollary 4.3-(a)]) that if one considers two connected graphs G and G' with the same vertex sets, but where the set of edges of G contains the set E' of edges of G', then

it is possible to couple the UST in G with the UST in G' in such a way that $\mathcal{T}(G) \cap E' \subset \mathcal{T}(G')$ almost surely.

Suppose now that $\mathcal{I}$ is a collection of edges of a finite graph G, and let $\mathcal{I}_1 \subset \mathcal{I}_2$ be two subsets of $\mathcal{I}$ that can be both completed into spanning trees of G by adding edges that are not in $\mathcal{I}$. Let $\mathcal{T}_1$ (resp. $\mathcal{T}_2$), be the uniform spanning tree $\mathcal{T}(G)$ on G, conditioned on $\mathcal{T} \cap \mathcal{I} = \mathcal{I}_1$ (resp. on $\mathcal{T} \cap \mathcal{I} = \mathcal{I}_2$). It is then possible to couple $\mathcal{T}_1$ and $\mathcal{T}_2$ in such a way that $\mathcal{T}_2 \cap \mathcal{I}^c \subset \mathcal{T}_1 \cap \mathcal{I}^c$ almost surely.

Indeed, one can first condition both USTs to contain all edges in $\mathcal{I}_1$ and no edge in $\mathcal{I} \setminus \mathcal{I}_2$ (and this corresponds to just removing the edges of $\mathcal{I} \setminus \mathcal{I}_2$ from the graph and to collapse all edges of $\mathcal{I}_1$). Hence, one needs only to treat the case where $\mathcal{I}_1$ is empty and $\mathcal{I}_2 = \mathcal{I}$, which can be deduced from the previously mentioned result by conditioning on $\mathcal{I} \cap \mathcal{T}_1$.

Again, these results have fairly obvious generalizations to the case of weighted graphs (we safely leave their proofs to the readers).

2.2. **Schramm's framework.** In order to describe the scaling limits of our forests, we will use the framework introduced by Oded Schramm [22]; let us briefly review its basic features (we refer to Section 10 of [22] for details).

For a compact topological space X, let us call $\mathcal{H}(X)$ the set of compact subsets of X equipped with the Hausdorff topology; recall that $\mathcal{H}(X)$ is itself a compact space.

We call *Schramm space* $\mathcal{OS}$ in the Riemann sphere $\hat{\mathbb{C}}$ the set $\mathcal{H}(\hat{\mathbb{C}} \times \hat{\mathbb{C}} \times \mathcal{H}(\hat{\mathbb{C}}))$ equipped with its Hausdorff topology. Similarly, when Ω is a simply-connected bounded domain of the plane with $\mathcal{C}^1$ boundary, we define $\mathcal{OS}_\Omega = \mathcal{H}(\overline{\Omega} \times \overline{\Omega} \times \mathcal{H}(\overline{\Omega})) \subset \mathcal{OS}$. The distance of this Hausdorff topology on $\mathcal{OS}$ or $\mathcal{OS}_\Omega$ is denoted by $d_{\mathcal{H}}$. Note that as we took Ω to be bounded, it does not matter whether one works initially with the spherical or Euclidean distance in Ω. All these spaces are compact, so that any sequence of probability measures on those spaces possesses subsequential limits.

In the framework of uniform spanning trees and their scaling limits, one considers very special elements in $\mathcal{OS}$ (in particular elements $\mathcal{G}$ in $\mathcal{OS}$ with the property that if $(a, b, K) \in \mathcal{G}$, then K is a continuous path from a to b, and $(b, a, K) \in \mathcal{G}$). A discrete graph embedded in the plane can be encoded by its path ensemble, i.e. by a point $\mathcal{G}$ in the Schramm space such that $\mathcal{G} = \bigcup (a, b, \gamma)$ where a, b run over all couple of points in the graph and γ runs over all simple (continuous) paths joining a to b in the graph. In particular, when a or b does not belong to the graph or if a and b are in different connected components, then there is no triplet of the form (a, b, γ) in the corresponding path ensemble.

The UST on a discrete graph embedded in the plane can then be viewed as a probability measure on $\mathcal{OS}$, and by compactness, it has subsequential limits when one lets the mesh of the lattice go to zero. As we shall now recall, this subsequential limit is in fact a limit: From now on and until further notice, Ω will denote either the entire plane or a simply-connected bounded domain of the plane with $\mathcal{C}^1$ boundary. We set Ω^δ a simply connected discretization of it at mesh size δ of the same type as in [19] ("union of squares" domain, paragraph 2.1): Fix $\xi \in \Omega$ and consider A the simply connected subset of $\delta\mathbb{Z}^2 \cap \Omega$ containing ξ. The set Ω^δ is the union of the squares $\{a + x + \mathbf{i}\, y,\ |x| \leq 1/2,\ |y| \leq 1/2\}$ over $a \in A$ (for instance, when Ω is the entire plane, just take Ω^δ to be $\delta\mathbb{Z}^2$). Let us consider the UST $\mathcal{T}(\Omega^\delta)$ and its path ensemble denoted by $\mathcal{G}^\delta(0) \in \mathcal{OS}_\Omega$.

The branches of uniform spanning trees are loop-erased random walks (LERW), which have been shown by Lawler, Schramm and Werner to converge to SLE_2 paths in the scaling limit (this convergence holds for paths parametrized by "Loewner capacity" which yields in particular convergence for paths up to monotone reparametrization), [17, Theorem 1.1].

As explained in [22], the convergence of LERW to SLE_2, together with estimates building on Wilson's algorithm yields the convergence of the UST to its continuous limit in the Schramm space (we will refer to results and statements that are proved in other papers or preprints as "results"

in order to make the distinction with the lemmas and propositions that are proved in the present paper):

Result A. ([17, Corollary 1.2] and [22, Theorem 11.3]). When $\delta \to 0$, $\mathcal{G}^\delta(0)$ converges in distribution (in $\mathcal{OS}_\Omega$) towards a continuous random element $\mathcal{G}(0)$.

There exist other possible descriptions of the scaling limits of USTs (for instance via the contour process of the tree, that converges to SLE_8) but we will not use them here. We will just call the random object $\mathcal{G}(0)$ the continuous UST in Ω. Theorem 1.5 of [22] lists various properties of $\mathcal{G}(0)$ (that for instance explain why one can call it a random tree). In particular, for every given $x, y \in \overline{\Omega}$, there exists almost surely a unique $\omega \in \mathcal{H}(\overline{\Omega})$ such that $(x, y, \omega) \in \mathcal{G}(0)$. Moreover, if $x \neq y$, then ω is almost surely a simple path, and if $x = y$, then ω is almost surely a single point. There are some random exceptional points, for which this uniqueness statement does not hold (these points are nonetheless well-understood, in terms of the dual tree). However, existence never fails i.e., almost surely, for any x and y, there exists at least one $\omega \in \mathcal{H}(\overline{\Omega})$ such that $(x, y, \omega) \in \mathcal{G}(0)$.

There are several ways to approximate the continuous UST in the Schramm space by somewhat simpler (continuous) objects. It is for instance natural to consider a dense deterministic sequence of points $z_1, z_2, \ldots$ in Ω and to define for each n the finite subtree $\mathcal{T}_{z_1,\ldots,z_n}$ consisting of just the branches that join $z_1, \ldots, z_n$ (we have seen that they are unique), and to see that when $n \to \infty$, this finite tree almost surely converges to the continuous UST in $\mathcal{OS}$. This last statement holds in a strong way: the finite trees approximate well the entire tree in the sense that for all $\varepsilon > 0$, the whole tree is formed of the finite tree $\mathcal{T}_{z_1,\ldots,z_{n_\varepsilon}}$ plus some paths of diameter smaller than ε with high probability. This key property was derived in [22] and we now state it more precisely.

In what follows, when Ω is the entire plane, we will use the spherical distance. We say that a subset $\mathcal{G}_\varepsilon$ of some $\mathcal{G} \in \mathcal{OS}_\Omega$ is a strong ε-approximation of $\mathcal{G}$ if for any point $(a, b, \omega) \in \mathcal{G}$, we can find a_ε, b_ε, ω_a, ω_b and ω' such that $d(a, a_\varepsilon) \le \varepsilon$, $d(b, b_\varepsilon) \le \varepsilon$, $(a_\varepsilon, b_\varepsilon, \omega') \in \mathcal{G}_\varepsilon$, $\omega_a \subseteq B(a, \varepsilon)$, $\omega_b \subseteq B(b, \varepsilon)$ and $\omega = \omega_a \cup \omega' \cup \omega_b$. When $\mathcal{G}$ encodes the branches of a tree, approximations of this kind can be found by somehow removing the part of the branch $(a, b, \gamma_{a,b})$ in an ε-neighborhood of a and b (in [22], Schramm defined the ε-trunk as a subtree of the UST where the part of the branches which are ε close to the leaves are removed. It is then obvious that the ε-trunk considered in the Schramm space is a strong approximation of $\mathcal{G}(0)$). Note that, in particular the distance between $\mathcal{G}_\varepsilon$ and $\mathcal{G}$ is smaller than ε. The following result is a key step in [22] towards the proof of Result A.

Result B. [22, Theorem 10.2] For any cut-off $\varepsilon > 0$, we can find a scale $\delta_\varepsilon > 0$, such that for any mesh size $\delta < \delta_\varepsilon$ and for all set of vertices $(z_1, \cdots, z_n)$ being a δ_ε-net of Ω (i.e. every point in Ω is within distance δ_ε of one of the z_i), the finite tree $\mathcal{T}_{z_1,\ldots,z_n}(\Omega^\delta)$ generated by $z_1, \cdots, z_n$ viewed in the Schramm space $\mathcal{OS}_\Omega$ is a strong ε-approximation of $\mathcal{G}^\delta(0)$ with probability greater than $1 - \varepsilon$.

As a consequence (see [22, Corollary 10.3]), for all $\varepsilon > 0$, there exists n such that the continuous subtree $\mathcal{T}_{z_1,\ldots,z_n}(\Omega)$ is a strong ε-approximation of the continuous UST $\mathcal{G}(0)$, with probability greater than $1 - \varepsilon$.

Dual trees and boundary conditions. It is well known that for a planar graph (i.e. embedded in the plane so that no two edges cross), one can associate to each spanning tree T on the graph G a dual spanning tree on the dual graph, and that if $\mathcal{T}$ is sampled according to the UST measure, then the dual tree is sampled according to the UST measure in the dual graph. When G is a portion of the lattice $\mathbb{Z}^2$, then the dual graph is a portion of the lattice $(\mathbb{Z} + 1/2)^2$, with the boundary vertices identified (this corresponds to wired boundary conditions). In the discrete case, one can define $\mathcal{G}^\delta(0)^\star$ in the Schramm space as being the dual tree of $\mathcal{G}^\delta(0)$ (i.e. the element in the Schramm space corresponding to the dual of the tree $\mathcal{T}(\Omega^\delta)$). By taking subsequential limits, one can then have convergence in distribution of the couple $(\mathcal{G}^\delta(0), \mathcal{G}^\delta(0)^\star)$. It is explained in [22] that in fact,

the limit of $\mathcal{G}^\delta(0)^\star$ is a deterministic function of the limit of $\mathcal{G}^\delta(0)$. We again refer to [22] for details (in particular about boundary conditions for the USTs).

Building on Wilson's algorithm, it is fairly easy to compare UST's with different boundary conditions, and to deduce the convergence (when the mesh size goes to 0) of the UST in the entire plane from the convergence in bounded domains. For instance, if one considers n points $y_1, \dots, y_n$ in the plane, and the law of the finite tree $\mathcal{T}_n^\delta$ obtained by sampling the (smallest) part of the UST in $\delta\mathbb{Z}^2$ that contains n points on this grid that are at distance smaller than δ from $y_1, \dots, y_n$, then the law of this tree will converge as $\delta \to 0$ to the law of a finite SLE_2-tree joining $y_1, \dots, y_n$. Furthermore, the law of this finite continuous tree is the limit when $R \to \infty$ of the law of the corresponding tree, in the domain $\{z \,:\, |z| < R\}$.

2.3. **UST and lengths of branches.** We now want to extend the previous convergence in distribution of the discrete UST to the continuous one, when one adds also the information about the lengths of the branches of tree. It is known since Rick Kenyon's paper [13] that the mean number of steps of a LERW grows like $\delta^{-5/4+o(1)}$ as the mesh-size δ goes to 0 (see also [21, 4] for closely related sharper estimates and results).

On the other hand, it is also known that the scaling limit of LERW (i.e. SLE_2) is a random simple curve with Hausdorff dimension 5/4 [3]. In fact, it has been recently shown [16] than SLE_2 can be naturally parametrized (i.e. as a continuous curve) by its 5/4-dimensional Minkowski content. Recall that the d-dimensional Minkowski content of a curve γ is defined as:

$$\mathrm{Cont}_d(\gamma) = \lim_{\varepsilon \to 0} \varepsilon^{d-2} \mathrm{Area}\{z \,:\, d(z,\gamma) \le \varepsilon\}$$

provided that the limit exists.

It is natural to expect that in fact, the suitably renormalized discrete length of the LERW should converge to the 5/4-dimensional content of the limiting SLE_2. This non-trivial fact turns out to be correct: Let Ω be a bounded simply connected domain with analytic boundary such that $0 \in \Omega$ and for each δ, recall that Ω^δ is a lattice approximation of Ω in $\delta\mathbb{Z}^2$. Consider a loop erased random walk starting at 0 in Ω^δ (i.e. the loop erasure of a simple random walk stopped when it hits $\partial\Omega^\delta$), that we view as a continuous curve that takes one unit of time to cross an edge, and denote by γ^δ its time-reversal. The following consequence of the main result of [19] will be an essential building block in our paper, that enables us to fine-tune the scale and control the cutting procedure.

Result C (Proposition 10 in Appendix A)**.** The curve $t \mapsto \gamma^\delta(t/\delta^{5/4})$ converges in distribution towards the radial SLE_2 in Ω (starting from a point chosen with respect to the harmonic measure on $\partial\Omega$ seen from 0) in its natural parametrization, for the topology of supremum norm (in order to accommodate the fact that the paths have different time-duration, one can let them stay at the origin when they hit it).

Combining Result C with Wilson's algorithm, the convergence of the discrete wired UST in bounded domains Ω yields readily the following more general results whose proofs are deferred in Appendix A:

- The convergence of the wired UST with arc-length parametrization (Corollary 11).
- The convergence of the plane UST (Proposition 12) and free UST (Proposition 13).
- The convergence of the joint law of the UST and its dual with their lengths thanks to Schramm result recalled in the previous paragraph 2.2.

2.4. **Scaling limit of the cutting dynamics.** In the following, Ω is either the entire plane or a simply connected bounded domain with C^1 boundary, and Ω^δ denotes its discretization at mesh size δ.

Let us now define the discrete cutting procedure. Let $(-\tau_e)$ be a family of i.i.d random exponential times with mean $\delta^{-5/4}$, indexed by the set of (non-oriented) edges e of Ω^δ. We start at time

$t = 0$ with a UST $\mathcal{G}^\delta(0)$ on Ω^δ independent of the family (τ_e). For a fixed time $t < 0$, we define $\mathcal{G}^\delta(t) \subseteq \mathcal{G}^\delta(0)$ to be the spanning forest that is obtained from $\mathcal{G}^\delta(0)$ by removing all the edges e with $\tau_e \in (t, 0]$ (viewed in the Schramm space, we remove all the paths that go through at least one of these edges). This defines a nested family of forests $(\mathcal{G}^\delta(t))_{t\leq 0}$. Note that the limit point $\mathcal{G}^\delta(-\infty)$ is a graph without edges (encoded in the Schramm space by the point $\bigcup_{v\in\Omega^\delta}(v, v, \{v\})$).

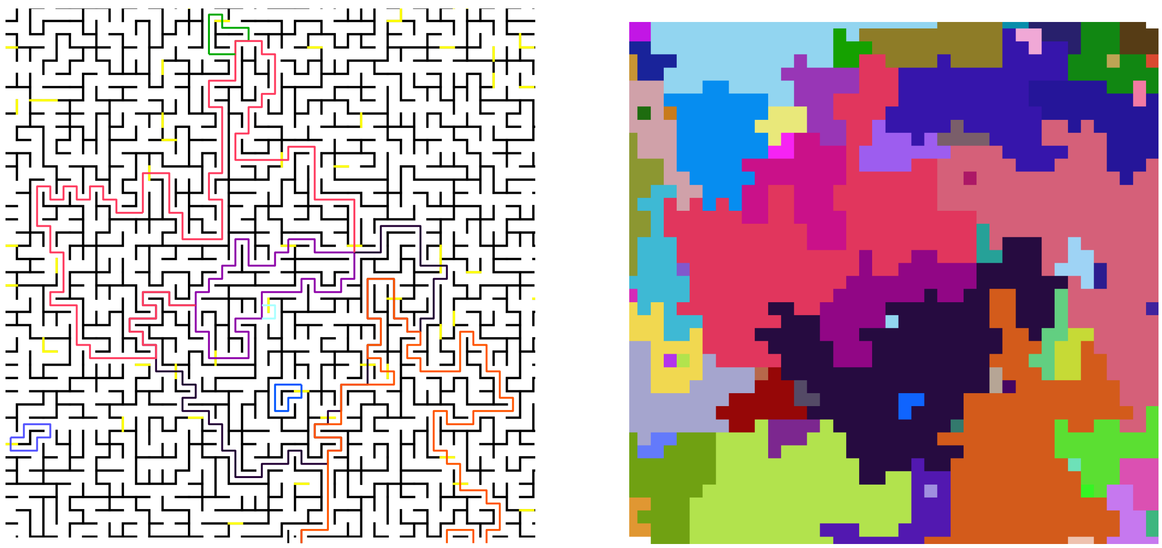

FIGURE 4. On the left hand side, the forest at time $t = -1$ of the dynamics with some of its clusters highlighted (the yellow edges were removed from the initial UST) and, on the right hand side, its connected components.

Let us now define the continuous counterpart of this discrete cutting procedure. We first sample (for a given Ω) the continuous UST $\mathcal{T} = \mathcal{G}(0)$. For any fixed $z_1, \dots, z_n$, the 5/4-Minkowski content of the tree $\mathcal{T}_{z_1,\dots,z_n}$ is finite. We then sample a Poisson point process on this finite tree where marked points appear with (negative) time with an intensity given by this 5/4-Minkowski content. As we do this simultaneously for any finite set of points z_i, we in fact are having marks appearing on the "backbone" of the continuous UST. We then define the continuous forest $\mathcal{G}(t)$ that corresponds to the continuous tree, by cutting all marked points that have appeared in the time-interval $(t, 0]$.

Note that when Ω is the entire plane, the underlying metric used to define the Schramm space is the spherical metric, but the cutting procedure uses the 5/4-dimensional content associated to the Euclidean metric, as it should correspond to the limit of the discrete length of the LERW on the graph.

Proposition 1. *The process $(\mathcal{G}^\delta(t))_{t\leq 0}$ converges in distribution (in the sense of finite-dimensional distributions in $\mathcal{OS}_\Omega$) towards the process $(\mathcal{G}(t))_{t\leq 0}$.*

Note that it is possible (and quasi immediate) also to state a similar result for a stronger Skorokhod-type convergence on càdlàg processes.

Proof. We fix t_0, and $\varepsilon, \eta > 0$, and find n such that with probability greater than $1 - \eta$, the finite subtree $\mathcal{T}_n := \mathcal{T}_{z_1,\dots,z_n}$ generated by $z_1, \cdots, z_n$ is a strong ε-approximation of $\mathcal{G}(0)$ i.e. differs from it by appending pieces of paths of diameter less than ε (by a slight abuse of notation, $\mathcal{T}_n$ will represent the tree both as a union of branches and as a point in Schramm space). In particular, to understand the cut forest $\mathcal{G}(t)$ up to a distance smaller than ε, it is enough to look at what is happening inside $\mathcal{T}_n$. Let us denote $\mathcal{T}_n(t)$ the cutting of the tree $\mathcal{T}_n$, i.e. the graph $\mathcal{T}_n \cap \mathcal{G}(t) \in \mathcal{OS}_\Omega$.

We similarly define in the discrete setting the subgraph $\mathcal{T}^\delta_n := \mathcal{T}^\delta_{z_1,\dots,z_n}$ of $\mathcal{G}^\delta(0)$, and $\mathcal{T}^\delta_n(t)$.

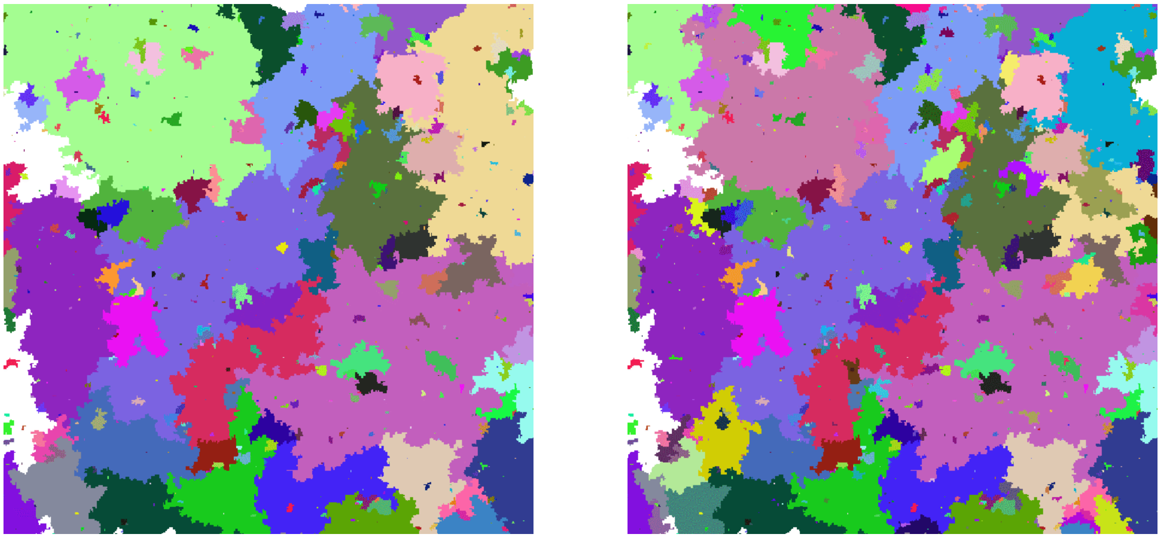

FIGURE 5. Simulations of $\mathcal{G}^\delta(3t)$ and $\mathcal{G}^\delta(4t)$ (different clusters are indicated in different colors): the latter is obtained from the former by cutting while the former is obtained by the latter via the glueing Markov process

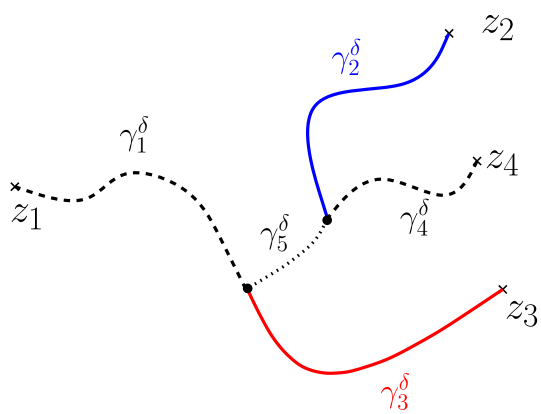

FIGURE 6. The simple paths $(\gamma_k^\delta, k = 1, \cdots, 5)$ of a representation of $\mathcal{T}_4^\delta$

Proposition 12 and 13 tell us that the finite subtrees $\mathcal{T}_n^\delta$, together with their length measure converge: any of the branches from z_i to z_j converges, in natural parametrization for the topology of supremum norm, towards branches of the continuous tree $\mathcal{G}(0)$.

The branches of the tree $\mathcal{T}_n^\delta$ can be divided into several disjoint paths. We choose to divide it according to the intersection points of its branches (see Fig. 6). Thus, $\mathcal{T}_n^\delta$ is seen as the reunion of N (at most and typically equal to $2n-3$) simple paths γ_k^δ. When the mesh size go to 0, there is a correspondence between the discrete paths γ_k^δ and the continuous ones (even if this can be checked to a.s. not happen, we allow for constant paths here). From Result C for the tree, the finitely many γ_k^δ converge in their natural parametrization towards their corresponding paths in the continuous tree for the topology of uniform convergence.

We couple the cutting dynamics in the discrete and in the continuum in the following way: assume the mesh size is small enough so that the lengths α_k^δ of γ_k^δ are ε' (to be fixed later) close to the lengths α_k of γ_k.

We then sample independent Poisson point processes of parameter $|t_0|$ on N intervals of respective lengths $\alpha_k + \varepsilon'$, and transfer these Poisson point processes on the discrete and continuous branches using respectively the parametrization by length and the natural parametrization (e.g. in the discrete setting, when a point falls in an interval of the type $[m\delta^{5/4}, (m+1)\delta^{5/4})$, we remove the corresponding edge). We now choose ε' such that with probability at least $1-\eta$, no point is drawn in any of the intervals $[\alpha_k - \varepsilon', \alpha_k + \varepsilon']$ and we work under this assumption: this allows us to be sure that there is a correspondence between discrete and continuous cut points. Note that this procedure gives a coupling of the dynamics at all time $t \in [t_0, 0]$, by associating to each point of the Poisson point processes independent uniform time labels in $[t_0, 0]$.

We also sample cutting points that lies outside of $\mathcal{T}_n^\delta$ (resp. $\mathcal{T}_n$). Nonetheless, however we have cut $\mathcal{T}_n$, cutting away any more points outside of it can displace it in Schramm space by a distance at most ε, as $\mathcal{T}_n$ is a strong ε-approximation of $\mathcal{G}(0)$.

Now choose the mesh size δ small enough such that with probability at least $1-\eta$, $\mathcal{T}_n^\delta$ and $\mathcal{T}_n$ are at distance ε in the distance of uniform convergence of pieces γ_k in their natural parametrizations. Therefore, with probability greater than $1-\eta$, for all time $t \in [t_0, 0]$, the cutting points of the previous coupling are ε-close and the Hausdorff distance between $\mathcal{T}_n^\delta(t)$ and $\mathcal{T}_n(t)$ is smaller than ε.

Gathering the previous estimates, on a set of probability $1-3\eta$, for any time $t \in [t_0, 0]$, the graphs $\mathcal{G}^\delta(t)$ and $\mathcal{G}(t)$ are at a distance at most $d_{\mathcal{H}}(\mathcal{G}^\delta(t), \mathcal{T}_n^\delta(t)) + d_{\mathcal{H}}(\mathcal{T}_n^\delta(t), \mathcal{T}_n(t)) + d_{\mathcal{H}}(\mathcal{T}_n(t), \mathcal{G}(t)) \leq 3\varepsilon$, which implies the convergence of the finite-dimensional marginals. □

3. The structure graph and the scaling limit of the glueing dynamics

Let us now focus on the flow that one obtains when one looks at the time-reversal of the cutting dynamics on some interval $[t, 0]$.

Description of the discrete glueing dynamics. Recall that if we are observing $\mathcal{G}^\delta(t)$ for some given $t < 0$, we can recreate the conditional law of $(\mathcal{G}^\delta(s))_{s\in[t,0]}$ in the following way: Denote by n the number of connected components of $\mathcal{G}^\delta(t)$. Let us pick uniformly a set of edges E among sets E' of n edges of $\delta\mathbb{Z}^2$ such that $\mathcal{G}^\delta(t) \cup E'$ is a spanning tree of Ω^δ. The graph $\mathcal{G}^\delta(t)$ then evolves by iteratively gaining edges of E (picked in uniform order), at the jump times of a Poisson process conditioned on jumping n times in $[t, 0]$ (or equivalently, edges of E appear at independent uniformly chosen times).

Let us rephrase this evolution in a way that is more tractable in the continuum limit. We first (deterministically) associate to each $\mathcal{G}^\delta(t)$ a structure graph $\mathcal{S}^\delta(t)$ as described in the introduction: Each connected component c of $\mathcal{G}^\delta(t)$ becomes a site of the structure graph $\mathcal{S}^\delta(t)$. Two neighboring (and distinct) connected components are linked by an edge in the structure graph, that carries a positive weight equal to $\delta^{5/4}$ times the number of edges in Ω^δ between the two connected components (edges with one end-point in each of the connected components).

It turns out that the trace of the set of edges E on the structure graph (which shows how the connected components of $\mathcal{G}^\delta(t)$ are connected in the graph $\mathcal{G}^\delta(0)$) has the law of the weighted spanning tree $\mathcal{T}(\mathcal{S}^\delta(t))$ on $\mathcal{S}^\delta(t)$. This describes the Markovian evolution of the discrete glueing dynamics when seen on structure graphs (each edge that is in the weighted tree then appear uniformly at random in the interval $[t, 0]$).

Note that the conditional distribution of the evolution of $(\mathcal{G}^\delta(s))_{s\in[t,0]}$ given the initial data $\mathcal{G}^\delta(t)$ and the evolution of the structure graph $(\mathcal{S}^\delta(s))_{s\in[t,0]}$ is easy to describe. When two sites c and c' of $\mathcal{S}^\delta(s)$ merge, then one chooses uniformly among the edges (in the discrete lattice picture) that join c and c' which one this merging corresponds to in the original graph.

Definition of the continuous structure graphs. The first non-trivial job when trying to make sense of the continuous counterpart of this glueing dynamics on structure graphs is to construct the continuous structure graphs $\mathcal{S}(t)$: Obviously, the vertices of $\mathcal{S}(t)$ shall be the connected components of $\mathcal{G}(t)$

and edges of $\mathcal{S}(t)$ link two vertices (i.e. connected components) whose corresponding components have a common boundary. The candidate for the weight of these edges is (up to a constant) the 5/4-dimensional Minkowski content of the interface between the corresponding clusters. Here we can note that this interface is in fact made of portions of branches in the dual tree, which suggests that we will need to control the lengths of the branches the dual of the continuous tree. This is the purpose of the next result (we defer its proof to Section 4) that then defines, for each $t \leq 0$, the weights of the structure graph $\mathcal{S}(t)$ and shows that they are indeed the limits of their discrete counterparts:

Proposition 2 (Weights of the continuous structure graph)**.** *Consider two given points z_0 and z_1 in Ω, and the connected components $c_0^\delta(t)$ (resp. $c_0(t)$) and $c_1^\delta(t)$ (resp. $c_1(t)$) of $\mathcal{G}^\delta(t)$ (resp. $\mathcal{G}(t)$) that they are part of, and let $l^\delta(z_0, z_1)$ be the renormalized length of the interface between $c_0^\delta(t)$ and $c_1^\delta(t)$ (respectively the 5/4-dimensional Minkowski content $l(z_0, z_1)$ of the intersection between $c_0(t)$ and $c_1(t)$) when it exists. Then, for each given t, the couple $(\mathcal{G}^\delta(t), l^\delta(z_0, z_1))$ converges in distribution to $(\mathcal{G}(t), l(z_0, z_1))$.*

Mind that this is not a trivial fact, because the structure graphs are rather complicated: we have to handle the infinitely many microscopic clusters appearing in the scaling limit and that will squeeze in between two macroscopic ones. One point in the proof will be to control the effect of this feature.

In order to define the Markov dynamics on such structure graphs, we will need to define the (weighted) forests and trees on them. In order to do so, we will choose exhaustions $(\mathcal{S}_\varepsilon(t))_\varepsilon$ and $(\mathcal{S}_\varepsilon^\delta(t))_\varepsilon$ of the graphs $\mathcal{S}(t)$ and $\mathcal{S}^\delta(t)$. Recall that the limiting laws on forests (when $\varepsilon \to 0$) do not depend on the choice for the exhaustions (see e.g. §5 of [6]) so that we are free to choose one that is well-tailored for our purposes (and we can also use geometric information from $\mathcal{G}(t)$ about the size of the clusters that correspond to the vertices of the structure graphs): For all $\varepsilon > 0$, we define the vertex set of $\mathcal{S}_\varepsilon^\delta(t)$ (resp. $\mathcal{S}_\varepsilon(t)$) to be the subset of the vertex set of $\mathcal{S}^\delta(t)$ (resp. $\mathcal{S}(t)$) consisting of the connected components $\mathcal{G}^\delta(t)$ (resp. $\mathcal{G}(t)$) that have a diameter at least ε (when Ω is the entire plane, we use the spherical metric here). The weighted edges between vertices of $\mathcal{S}_\varepsilon^\delta(t)$ and $\mathcal{S}_\varepsilon(t)$ are then exactly those of $\mathcal{S}^\delta(t)$ and $\mathcal{S}(t)$. From Result B, it follows that the graphs $\mathcal{S}_\varepsilon^\delta(t)$ and $\mathcal{S}_\varepsilon(t)$ are almost surely finite. It is also immediate to see that $(\mathcal{S}_\varepsilon^\delta(t))_\varepsilon$ (resp. $(\mathcal{S}_\varepsilon(t))_\varepsilon$) exhausts $\mathcal{S}^\delta(t)$ (resp. $(\mathcal{S}(t))$.

We now state the convergence of the structure graph. We use the discrete topology on finite graphs, and for a given finite graph, weights form a real vector space that we equip with its natural topology.

Corollary 3 (Discrete to continuous structure graph convergence)**.** *For each $t < 0$, for all but (at most) countably many positive ε, the finite random graph $\mathcal{S}_\varepsilon^\delta(t)$ converges in probability to $\mathcal{S}_\varepsilon(t)$ as the mesh size δ goes to 0.*

This results follows directly from Proposition 2 (i.e. the convergence of the weights of the edges) and the convergence of $\mathcal{G}^\delta(t)$ to $\mathcal{G}(t)$. The presence of the constraint on ε is just to ensure that for those values of ε, almost surely no diameter of cluster in $\mathcal{G}(t)$ is exactly equal to ε (as this would potentially create a problem in the limit of the cut-off). As we know that there are countably many clusters, it follows that this bad scenario can anyway happen for at most countably many ε (for each fixed t). We could of course also (try to) prove that this anyway never happens, but the present result will be enough for our purposes. The proofs of those technical results are deferred to Section 4.

Abstract definition of the Markovian dynamics on structure graphs. We are now ready to define the Markovian dynamics on the set of structure graphs. For a given t and a given weighted graph $\mathcal{S}(t)$:

- First, sample a weighted free spanning forest on $\mathcal{S}(t)$, and for each edge of this forest, sample independently a uniform random variable on $[t,0]$ that indicate when this edges appears.
- Then, construct the graph at time $s \in [t,0]$ by contracting all edges that have appeared before time s, and using the addition rule for weights (loosely speaking, when two sites s_1 and s_2 merge into a site s_1s_2, the new weights are given by $w_{new}(s_1s_2,\cdot) = w_{old}(s_1,\cdot)+w_{old}(s_2,\cdot)$.

Recall that it is not a priori clear that the weighted spanning forest on the structure graph is a tree, but along our proof, we will see that in fact, it is indeed almost surely the case, when one starts this dynamics with the random graph $\mathcal{S}(t)$. Moreover, weights can blow up under the dynamics, depending on initial conditions. That this does not happen when we initiate our dynamics with the structure graphs of our near-critical spanning forests is a consequence of the following Theorem 4.

In this way, one defines a process $(\tilde{\mathcal{S}}(s))_{s\in[t,0]}$ which is the evolution of this Markovian dynamics when applied to the random structure graph $\tilde{\mathcal{S}}(t) = \mathcal{S}(t)$. The core of the matter is then to prove the following fact:

Theorem 4. *The law of $(\tilde{\mathcal{S}}(s))_{s\in[t,0]}$ is the same as that of $(\mathcal{S}(s))_{s\in[t,0]}$.*

In loose words, the scaling limit of the Markov dynamics on discrete structure graphs is Markov, and it is described by the simple process on continuous graphs that we have described above. Mind that the theorem is also valid when Ω is the full plane.

Note that, as in the discrete case, there is a (heuristically straightforward) description of the conditional distribution of $(\mathcal{G}(s))_{s\in[t,0]}$ given $\mathcal{G}(t)$. Construct first $\mathcal{S}(t)$ and $(\tilde{\mathcal{S}}(s))_{s\in[t,0]}$. For each concatenation of vertices $s(c)$ and $s(c')$ happening on $[t,0]$, we choose a point w according to the uniform measure on the common boundary of c and c', measured by its 5/4-dimensional Minkowski content (this common boundary is the union of several portions of dual branches, and its content is well defined, as follows from Lemma 8). Let us call $\mathcal{W}(s)$ the countable set of points thus chosen that corresponds to concatenations happening before time s. For each integer n, let $\tilde{\mathcal{G}}_n(s)$ be the reunion of the paths (a,b,γ), such that γ is a path from a to b that can be realized as the concatenation of at most n paths in $\mathcal{G}(t)$, where the points of concatenation belongs to the set $\mathcal{W}(s)$. We then define $\tilde{\mathcal{G}}(s)$ to be the closure in $\mathcal{OS}$ of the reunion $\cup_n \tilde{\mathcal{G}}_n(s)$. It is easy to see that $(\tilde{\mathcal{G}}(s))_{s\in[t,0]}$ has the same law as the limit of the discrete dynamics $(\mathcal{G}(s))_{s\in[t,0]}$. Indeed each given branch $(a,b,\gamma_{a,b}) \in \mathcal{G}(0)$ is almost surely cut a finite number of times, and there almost surely exist a countable family of branches of $\mathcal{G}(0)$ that are dense among the set of all branches of $\mathcal{G}(0)$ (see Result B).

Let us now explain how to deduce this theorem from the previous propositions. As we shall see, this is quite a soft argument, where we will exploit the tightness-type properties of the USTs (derived by Schramm) and coupling ideas.

Proof of Theorem 4. Let us first recollect a few facts:

1. From Result B, we know that for a given η and a given ε, we can find a finite set of points $z_1,\dots,z_n$, such that (for both the discrete case for all given δ, and the continuous case), with probability at least $1-\eta$, the graphs obtained by just cutting the trees $\mathcal{T}^\delta_{z_1,\dots,z_n}$ and $\mathcal{T}_{z_1,\dots,z_n}$ in a Poissonian way along its branches, do contain respectively the vertices and edges of the graphs $\mathcal{S}^\delta_\varepsilon(t)$ and $\mathcal{S}_\varepsilon(t)$.
2. On the other hand, for a given choice of $z_1,\dots,z_n$, the convergence of the branches of the tree joining these points in their natural parametrizations ensures that one can find ε_1 small enough so that (uniformly in δ i.e. for each given δ) the graph obtained by just cutting the finite trees $\mathcal{T}^\delta_{z_1,\dots,z_n}$ and $\mathcal{T}_{z_1,\dots,z_n}$ in a Poissonian way have the property that they are included in $\mathcal{S}^\delta_{\varepsilon_1}(t)$ and $\mathcal{S}_{\varepsilon_1}(t)$ with probability at least $1-\eta$ (simply because the probability that two cuts out of the finitely many cuts end up being at distance smaller than ε_1 of each other is very small).

(3) By the comparison results recalled at the end of Subsection 2.1, the law of the weighted spanning forest in $\mathcal{S}^\delta(t)$ when restricted to the edges in $\mathcal{S}^\delta_{\varepsilon_1}(t)$ is dominated by the law of the weighted spanning forest in $\mathcal{S}^\delta_{\varepsilon_1}(t)$, and the law of the weighted spanning forest in $\mathcal{S}(t)$ when restricted to the edge in $\mathcal{S}_{\varepsilon_1}(t)$ is dominated by the law of the weighted spanning forest in $\mathcal{S}_{\varepsilon_1}(t)$. In particular, if we are given n sites $s_1, \ldots, s_n$ and see that the tree in the weighted spanning forest in $\mathcal{S}^\delta(t)$ that joins these n points does stay in the graph $\mathcal{S}^\delta_{\varepsilon_1}(t)$ with probability at least A, then this means that one can couple the weighted spanning forest in $\mathcal{S}^\delta(t)$ and $\mathcal{S}^\delta_{\varepsilon_1}(t)$ in such a way that these two subtrees coincide with probability at least A (and the similar statement holds without the superscript δ).

(4) Finally, from Corollary 3, we know that for a well-chosen ε_1, the law of the weighted spanning forest on $\mathcal{S}^\delta_{\varepsilon_1}(t)$ converges to that of the weighted spanning forest on $\mathcal{S}_{\varepsilon_1}(t)$ as $\delta \to 0$.

Recall that $(\tilde{\mathcal{S}}(s))_{s\in[t,0]}$ is reconstructed from $\mathcal{S}(t)$ by sampling a weighted spanning forest on $\mathcal{S}(t)$ i.e. the limit of a weighted spanning forest in $\mathcal{S}_\varepsilon(t)$ as $\varepsilon \to 0$. On the other hand, $(\mathcal{S}(s))_{s\in[t,0]}$ is reconstructed by taking the limit when $\delta \to 0$ of the weighted spanning forest on $\mathcal{S}^\delta(t)$ (indeed, one reconstructs first $\mathcal{S}^\delta(s)$ and then takes the limit $\delta \to 0$).

Combining 1. and 2. shows that for all ε, one can find ε_1 small enough such that for all given δ, the subgraphs of $\mathcal{T}(\mathcal{S}(t))$ and of $\mathcal{T}(\mathcal{S}^\delta(t))$ that join all the sites of $\mathcal{S}_\varepsilon(t)$ and $\mathcal{S}^\delta_\varepsilon(t)$ stay respectively in $\mathcal{S}_{\varepsilon_1}(t)$ and $\mathcal{S}^\delta_{\varepsilon_1}(t)$ with probability greater than $1-2\eta$. By 3., we see that it is therefore possible to couple these subgraphs with those obtained when sampling $\mathcal{T}(\mathcal{S}_{\varepsilon_1}(t))$ and $\mathcal{T}(\mathcal{S}^\delta_{\varepsilon_1}(t))$ instead of $\mathcal{T}(\mathcal{S}(t))$ and of $\mathcal{T}(\mathcal{S}^\delta(t))$ so that they actually coincide with probability greater than $1-2\eta$. But by 4., we know that for all δ small enough, these two samples can be coupled so to be very close, which concludes the proof. Note that the argument also shows that the free spanning forest $\mathcal{T}(\mathcal{S}(t))$ is a.s. connected, hence a tree.

Mind that the identity in law between the two processes means the identity in law of each finite-dimensional marginals. And for any $t < s_1 < \ldots < s_n < 0$, we can always choose all the ε's and ε_1's in the above argument among those for which the convergence in Corollary 3 holds for these times $t, s_1, \ldots, s_n$.

Whole plane dynamics and its properties. Let us first observe that the previous Markov chain on structure graphs was not time-homogeneous. It was defined for all $t < 0$, on the time-horizon $[t,0]$ (i.e. for a time $|t|$) as follows: First sample the USF on the structure graph, and then open each edge e of this USF independently, at a uniformly chosen time $\tau(e)$ in $[t,0]$ independently.

However, it is trivial to turn this into a time-homogeneous Markov chain. One just needs to replace the uniformly chosen times in $[t,0]$ by (positive) exponential random variables $\xi(e)$ with mean 1 (one exponential variable for each edge of the structure graph), i.e. we do the time change $\xi(e) = \log(t/\tau(e))$. Then, the edge e opens at time $\xi(e)$ and one collapses it to form a new structure graph. As we shall now try to point out, this homogeneous-time Markov chain for structure graphs set-up turns out to be particularly interesting in the whole-plane setting.

Let us summarize the construction of the cutting dynamics $(\mathcal{G}(t))_{t\le 0}$ in the plane: Sample a continuous UST in the entire plane, and just as in the finite-volume case, define a Poisson point process on its branches, with intensity $\ell \times \mu$ where ℓ is the Lebesgue measure on $(-\infty, 0]$ and μ is the 5/4-dimensional Minkowski content measure. Then, for each $t < 0$, one can cut the UST on these marked points as before, which gives rise to a collection of trees $\mathcal{G}(t)$, and these trees are the limit when $\delta \to 0$ of their discrete counterparts $\mathcal{G}^\delta(t)$. Note that each site c of the structure graph $\mathcal{S}(t)$ is a cluster i.e. a subset of the plane. We know that the law of the continuous UST in the whole plane is scale-invariant. It therefore immediately follows that the processes $(\mathcal{G}(t))_{t\le 0}$ and $(\mathcal{S}(t))_{t\le 0}$ are scale-invariant too, in the following sense: For each $\lambda > 0$, we define $U_\lambda(\mathcal{G}(t))$ to be the set obtained by magnifying all clusters by a factor λ, and $U_\lambda(\mathcal{S}(t))$ the graph obtained by also

magnifying all edge-weights by a factor $\lambda^{5/4}$ (in other words, one relabels the edges of $\mathcal{S}(t)$ by just multiplying them by a factor $\lambda^{5/4}$. Then, the process $(U_\lambda(\mathcal{S}(t)))_{t\le 0}$ is identical in distribution to the process $(\mathcal{S}(t/\lambda^{5/4}))_{t\le 0}$.

Let us now define π to be the distribution of $\mathcal{S}(-1)$. Theorem 4 then states exactly that the process $(\mathcal{S}(-e^{-u}))_{u\ge 0}$ is obtained by letting the (time-homogeneous) Markov dynamics run from $\mathcal{S}(-1)$. But by the scale-invariance property, we get that (modulo relabeling of the edges of the structure graph), the distribution π is invariant under the time-homogeneous Markovian dynamic.

Finally, we can also note that if we start from the graph $S_0 = \mathbb{Z}^2$ with all edge-weights equal to 1 (or any other regular planar lattice) and let the time-homogeneous Markov chain $(S_u)_{u\ge 0}$ run until a large time U, we discover each edge of the final UST on $\mathbb{Z}^2$ (independently) with probability $1-e^{-U}$ (or more exactly, rather than their edges, their "traces on the structure graphs"). In particular, with Theorem 4, this shows that (modulo relabeling of the edges of the structure graph i.e. scaling down $\mathbb{Z}^2$ to $\delta\mathbb{Z}^2$ for an appropriately chosen δ depending on U), that as $U \to \infty$, the law of the structure graph converges to π (in the sense of Corollary 3).

Hence, this provides the following renormalization flow description of the UST scaling limit via (a rescaling of) the time-homogeneous Markov chain P_u on the state of discrete weighted graphs:

Theorem 5 (Renormalization flow description)**.** *The measure π (that describes the previous scaling limit of near critical spanning forests) is invariant under the Markov chain. Furthermore, the (time-homogeneous) Markov chain started from any deterministic periodic two-dimensional transitive lattice and properly rescaled does converge in distribution to π.*

4. Technical estimates and proofs

4.1. **First comments about the structure graphs and their convergence.** Most of the remainder of this paper is now devoted to the proof of Proposition 2, that provides the convergence of the discrete structure graph weights to their continuous counterpart. In this section, we are working with the UST on the whole plane but the proofs can easily be extended to any bounded domain with C^1 boundary.

Let us now make some comments about this, and explain how to deduce Proposition 2 from two lemmas that we will then prove in the subsequent section, based on more "traditional" arm-estimates and considerations for UST.

Suppose first that z_0 and z_1 are two given points. In both the discrete and continuous settings, these two points are joined by a unique path in the UST, that has a finite (renormalized) length (or Minkowski content – by slight abuse of terminology, we will now use the word length also in the continuous case), so that the number of "cuts" on this branch (conditional on this length, and for a given t) follows a Poisson distribution. If these two points z_0 and z_1 end up in different trees at the end of the cutting procedure, then there as been a "first cut" i.e. an edge e on this path that has been removed first (when one looks back from time 0 to time t in the cutting procedure), and its law (conditional on the branch between z_0 and z_1) is uniform on this branch with respect to length. Mind that the edge e has a positive probability not to exist (if there were no cut on the branch).

If we consider the entire UST and removes from it just this one edge e, then one has divided the UST into two trees, one containing z_0 and the other one containing z_1. The interface between these two trees is then described by a cycle $\mathcal{C}^\delta$ that consists of the edge $e^\star$ dual to e together with the branch in the dual of the UST, that joins (in the dual tree) the two extremities of $e^\star$. Clearly, if one removes more edges than just e, the trees that contain z_0 and z_1 respectively will shrink (it may become empty), and the interface between these two trees can only decrease. Hence, the interface between the two clusters of $\mathcal{G}^\delta(t)$ that contain z_0 and z_1 is a subset of this cycle (and its length is bounded by that of $\mathcal{C}^\delta$). The same situation occurs in the continuous case. Here, when one chooses a first point z at random (according to Minkowski-content) on the UST branch joining z_0 and z_1,

one can consider the cycle $\mathcal{C}$ in the dual tree that joins z to itself, and when one removes more points according to the cutting dynamics, the clusters that contain the two points z_0 and z_1 will intersect along a subset of that cycle $\mathcal{C}$. Note that we have a first easy lemma, that follows from

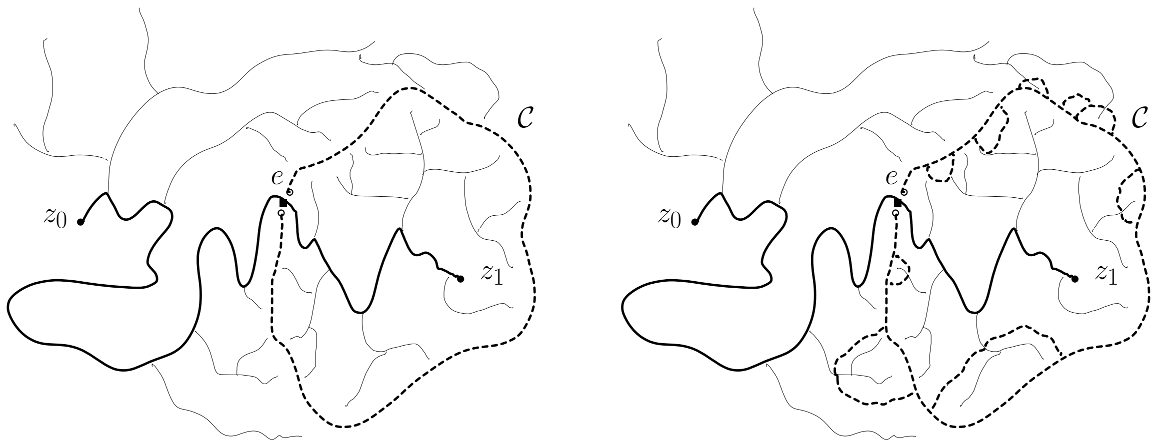

FIGURE 7. Sketch of the tree, of the cycle $\mathcal{C}$ and the cuts

the convergence in distribution of $\mathcal{C}^\delta$ and from the convergence of the renormalized measure on the branch from z_0 to z_1. In the following, $B(z, r)$ denotes the Euclidean ball of radius r centered in z when $z \in \mathbb{C}$ and when e is an edge of $\mathbb{C}^\delta$, $B(e, r)$ is the ball of radius r centered at its midpoint.

Lemma 6. *As $\eta_0 \to 0$, the probability that either $\mathcal{C}^\delta \not\subset B(0, 1/\eta_0)$, or $d(z_0, e) < \eta_0$ or $d(z_1, e) < \eta_0$ occurs goes to 0 uniformly with respect to δ.*

We know already that the lengths of branches in the dual tree converges to their continuous counterparts, but a little additional care will be needed when we want to deal with the length of the entire cycle $\mathcal{C}^\delta$, because it does originate at a special point i.e. a point on the backbone of the original UST, so we need to exclude the scenario where something weird happens to length of $\mathcal{C}^\delta$ in the vicinity of this special point. This is the purpose of the next lemma:

Lemma 7. *Let us fix η_0, z_0 and z_1, and condition on the event that $\mathcal{C}^\delta$ exists, and that the three events in Lemma 6 do not occur (note that this is a conditioning on an event of positive probability, bounded from below independently of δ, and that then, the diameter of $\mathcal{C}^\delta$ is bounded from below). As η goes to 0, in the previous setting (for fixed z_0 and z_1), the expected (conditional) renormalized length $u^\delta(\eta)$ of the two-sided part of $\mathcal{C}^\delta$ from $e^\star$ up to its first exits of the ball of radius η around the center of $e^\star$, does tend to 0 uniformly with respect to δ.*

Next, one can make the following observations (that can be made rigorous, but they serve here as a motivation and won't be used later, so we will not bother to do so): Suppose that in the previous scenario, one considers the continuous tree containing z_1 after cutting away just e, and that this tree is bounded (if we were in the whole plane, this means that z_1 was on the bounded side of the cut e). Lemma 7 shows that the length of $\mathcal{C}$ (in terms of Minkowski content) is finite, but one may wonder what the sub-tree containing z_1 does really look like at the end of the cutting procedure at time t, when one has removed from it many more cuts. One can notice that for a "typical point" on the cycle $\mathcal{C}$, the (Minkowski-content) length between this point z and z_1 in the initial tree is finite. Hence, it will have a positive probability to be cut off from z_1, but it also has a positive probability not to be cut off. Hence, the expected portion of the length of the part of $\mathcal{C}$ that will remain on the outer boundary of the cluster containing z_1 is in fact positive. On the other hand, a back-of-the envelope calculation (that we do not reproduce here) suggests that the total length of

the tree consisting of all the branches that join z_1 to all the boundary points in $\mathcal{C}$ is infinite. This means that an infinite number of macroscopic pieces of $\mathcal{C}$ are being cut out.

The purpose of the following lemma is now to control this feature at the discrete level: Let us say that a point z of $\mathcal{C}^\delta$ is cut-out from this boundary at a scale smaller than ε if there exists a cut disconnecting z from one of the two extremities of the special edge e, in such a way that the part of the tree disconnected from e by this cut has a diameter smaller than ε. For each $\eta > 0$, we are going to define L^δ to be the renormalized length of the set of points on $\mathcal{C}^\delta \cap (B(0, 1/\eta) \setminus B(e, \eta))$ that are cut-out from the interface $\mathcal{C}^\delta$ at a scale smaller than ε:

Lemma 8. *As ε goes to 0, in the previous setting (for fixed z_0, z_1 and η), the expected value of L^δ tends to 0 uniformly with respect to δ.*

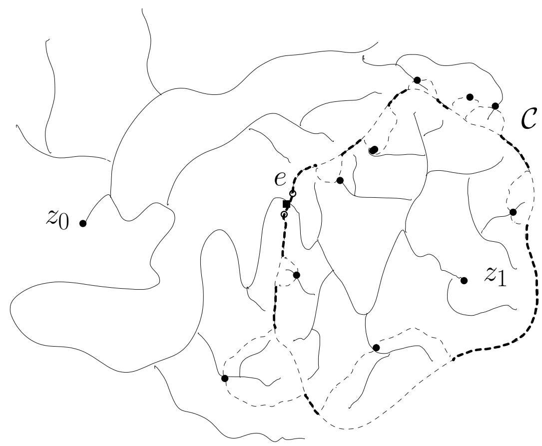

FIGURE 8. After all the cuts: The remaining interface between the trees containing z_0 and z_1

We shall prove Lemmas 7 and 8 in the next section, but let us already explain now how Proposition 2 follows from them:

Proof of Proposition 2. We fix z_0 and z_1. It suffices to control (for each η_0) the convergence of the weights on the event described in Lemma 6, i.e. when the interface is not too large, and when the cut e occurs neither near z_0 nor near z_1.

Suppose that we choose a sequence $\delta_k \to 0$ and graphs $\mathcal{G}^{\delta_k}(t)$ for all k, together with $\mathcal{G}(t)$ on the same probability space, in such a way that the tree containing (a δ_k-approximation of) z_0 (resp. z_1) in $\mathcal{G}^{\delta_k}(t)$ converges almost surely to the tree containing z_0 (resp. z_1) in $\mathcal{G}(t)$. We can furthermore assume that the dual tree at time 0 converges in such a way that the renormalized lengths of branches of its finite subtrees do.

The results by Schramm on strong approximations (for the dual tree), together with Lemma 7 ensures that the renormalized lengths of the discrete circuits $\mathcal{C}^{\delta_k}$ converge in probability to the 5/4-dimensional Minkowski content of their continuous counterpart $\mathcal{C}$ (and in fact the discrete circuit parametrized by renormalized length converges to the continuous one parametrized by Minkowski-content). Let us explain in a few words why this convergence holds, as even though we know the discrete cut points converge to the continuous ones, the dual cycles they close might a priori differ on a scale $\eta > 0$ in the following situation. Consider a point z of the primal tree where three distinct branches of diameter larger than η connect to each other. Then, if we consider two cut points close to such a branching point z, the dual cycles they close can fail to merge before exiting the ball $B(z, \eta)$, and as a consequence can widely differ. We actually already dealt with this issue in the proof of Proposition 1, when, for a fixed η, we provided (with high probability) a coupling of the

discrete and continuous cut points processes where they are chosen not only to be close, but also in a way that respect branching points of the tree at scale η.

Let us fix ε. Denote by l^δ_ε (resp. l_ε) the renormalized length of the set of points in $\mathcal{C}^\delta$ (resp. $\mathcal{C}$) that have not been disconnected from z_0 or z_1 at a scale larger than ε (i.e. by a cut creating a cycle of diameter larger than ε). In other words, we remove from the total length of $\mathcal{C}^\delta$ (resp. $\mathcal{C}$) the contribution of all the macroscopic cuts (of diameter larger than ε).

Now, the (finitely many) pieces of $\mathcal{C}^\delta$ cut by cycles of diameter larger than ε converge almost surely towards their continuous counterpart (for the same reason that $\mathcal{C}^\delta$ converges towards $\mathcal{C}$). In particular, we have that almost surely

$$\liminf_{k\to\infty} l^{\delta_k}_{2\varepsilon} \geq l_\varepsilon \geq \limsup_{k\to\infty} l^{\delta_k}_{\varepsilon/2}.$$

Note also that by definition, $l^\delta_\varepsilon \geq l^\delta(z_0,z_1)$. Hence, $l_\varepsilon \geq \limsup_{k\to\infty} l^{\delta_k}(z_0,z_1)$ almost surely.

Lemma 8 ensures on the other hand that $\mathbb{E}(l^{\delta_k}(z_0,z_1) - l^{\delta_k}_{2\varepsilon})$ goes to 0 as $\varepsilon \to 0$ uniformly in k. By definition of the continuous dynamics and of $l(z_0,z_1)$, we know that $l_\varepsilon \to l(z_0,z_1)$ almost surely as $\varepsilon \to 0$. It therefore follows from the previous inequalities that $l^{\delta_k}(z_0,z_1)$ converges in distribution towards $l(z_0,z_1)$. □

4.2. **Arm events in UST.** Let us first recall an estimate about LERW of the type that is essential in the derivation of results involving the Minkowski-content in [1, 4, 19]: Let X and Y be two independent simple random walks on $\mathbb{Z}^2$ starting at x and 0 respectively and stopped at their first exit time τ_X and τ_Y of the ball of radius N around the origin. Let us consider the loop erasure $\hat{Y}$ of Y. We denote by $\hat{Y}^L$ the subpath of $\hat{Y}$ from its last hitting time of the ball of radius L around the origin and define

$$\mathrm{Es}(L,N) := \mathbb{P}_{x=0}(X \cap \hat{Y}^L = \emptyset).$$

Result D. There exists a constant $C>0$ such that for all L and N with $L \leq N/2$,

$$C^{-1}(L/N)^{3/4} \leq \mathrm{Es}(L,N) \leq C(L/N)^{3/4}$$

Proof. When $L=1$, the estimate can be derived following the proof of [2, Corollary 3.15], using the better estimate of [4, Theorem 1.1] as an input. Moreover, one can compare $\mathrm{Es}(L,N)$ to $\mathrm{Es}(1,N)/\mathrm{Es}(1,L)$ thanks to [21, Propositions 5.2 and 5.3], which proves Result D. □

Note that this implies that the probabilities, say, $\mathrm{Es}(L,N)$ and $\mathrm{Es}(5L,N)$ are comparable. A further simple observation is that if we start a random walk anywhere in the disc of radius, say, $4L$ around the origin, the law of its hitting distribution of $\partial B(5L)$ is absolutely continuous (with Radon-Nikodym derivative bounded above and below, uniformly in L) with respect to the same hitting distribution when X starts from the origin. Hence, one has for instance, using the estimate for $\mathrm{Es}(5L,N)$, that

$$\mathbb{E}_Y\left(\sup_{x\in B(4L)} \mathbb{P}_x(X \cap \hat{Y}^{5L} = \emptyset)\right) \leq c\mathrm{Es}(5L,N) \leq c\,C\,5^{3/4}\,(L/N)^{3/4}, \tag{1}$$

which is a result that will be useful later on.

Note that the case where $L=1$ provides also directly (via Wilson's algorithm) the probability that two distinct branches in the wired UST in $B(N)$ that start at the origin and next to the origin do stay disjoint until they touch the circle of radius N. By duality, this is (almost, as there is the issue of the $(1/2,1/2)$ translation that we will not bother to mention in the following lines) exactly the probability that for the free UST in $B(N)$, there exists a branch from the boundary to the boundary that goes through a given vertex $x^\star$ to the origin.

An event related to the previous non-intersection events, but slightly different is the following arms event $\mathcal{A}(L, N)$ around the origin between scales L and N, that there exist four disjoint branches, two of the UST and two of the dual UST in alternating trigonometric order around the origin, that connect $\partial B(L)$ to $\partial B(N)$ (see Fig. 9).

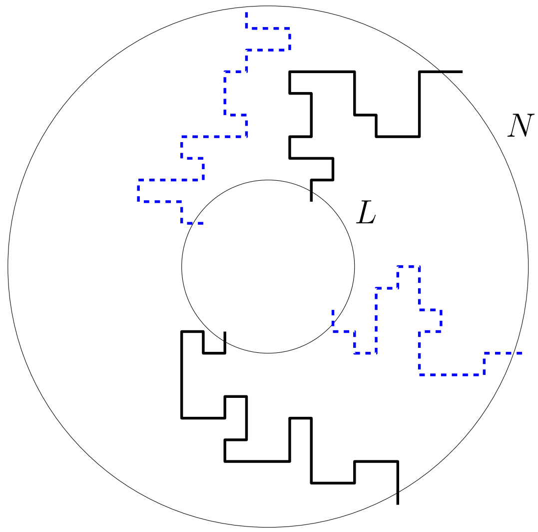

FIGURE 9. The four alternating disjoint branches in the UST.

Lemma 9. *Consider the UST in a discrete domain $\Omega \subseteq \mathbb{Z}^2$ containing $B(N)$, with arbitrary boundary conditions. There exists a constant $C > 0$ independent of Ω, N and the boundary conditions, such that for all $L \leq N/10$,*

$$\mathbb{P}(\mathcal{A}(L, N)) \leq C(L/N)^{3/4}.$$

The proof of this result will be based on the one hand on Result D, and on the other hand on the following simple estimate from [5]: Consider an UST $\mathcal{T}$ in a discrete domain $\Omega \subseteq \mathbb{Z}^2$ containing $B(3L)$, with arbitrary boundary conditions. Let us call K the maximal number of paths one can find in $\mathcal{T} \cap (B(3L) \setminus B(L))$ that touch both circles $\partial B(L)$ to $\partial B(3L)$ and that are not only disjoint but also disconnected in $B(3L) \setminus B(L)$.

Result E (Consequence of Theorem 2.2 in [5])**.** There exists a universal constant C (independent of L, Ω and the boundary conditions) such that $\mathbb{E}(K) < C$.

Proof of Lemma 9. Let us first explain why it suffices to prove the result in the case where $\Omega = B(N)$ with free boundary conditions. Note that (for whatever Ω and boundary conditions), when $\mathcal{A}(L, N)$ occurs, then one of the following two scenarios happens:

- Two branches of the UST that cross the annulus are joined in the UST inside of $B(L)$.
- Two branches of the dual UST that cross the annulus are joined in the UST inside of $B(L)$.

(it may happen that the two events occur simultaneously if there are more than four disjoint crossings). Hence, by symmetry and duality, we just have to evaluate the probability of the first event. But the stochastic coupling and monotonicity for USTs in different domains shows that the probability of this event is maximal (among all domains and boundary conditions) on the ball $B(N)$ with free boundary conditions on its boundary $\partial B(N)$.

By considering now the dual tree, we want to bound the probability that, for an UST in $B(N)$ with wired boundary conditions, there exist two disjoint branches of the tree that joint $\partial B(L)$ to

the outer wired boundary $\partial B(N)$. Since the branch γ of the tree from the origin to the boundary is always a branch that crosses the annulus, we see that we are actually are after the probability that there exists another branch in the UST $\mathcal{T}$ (with wired boundary conditions) that stays disjoint from γ and crosses the annulus from $\partial B(L)$ to $\partial B(N)$.

Let us look at the (dual) UST $\mathcal{T}$ in $B(N)$ conditioned on the branch γ as well as on all the edges of the UST inside $B(3L)$ (let us call $\mathcal{F}$ this σ-field generated by γ and the part of the UST in $B(3L)$). Conditionally on $\mathcal{F}$, the part of the UST in $B(N) \setminus (B(3L) \cup \gamma)$, is just a UST in this domain with wired boundary conditions on $\gamma \cup \partial B(N)$ and some more convoluted mixed boundary conditions on $\partial B(3L)$ (vertices on $\partial B(3L)$ are identified if they are connected by the UST inside of $B(3L)$). Let us denote by K the number of connected components of the intersection of the UST with $B(3L)$ that do touch both $\partial B(L)$ and $\partial B(3L)$. Recall from Result E that the expectation of K is bounded by an absolute constant. Let us arbitrarily choose points $x_1, \cdots, x_K \in \partial B(3L)$ in each of these connected components. The probability that $\mathcal{A}(L, N)$ holds is then bounded by the event that if we launch independent random walks (note that these walks will seem to jump around on $\partial B(3L)$, in a way that respect the boundary conditions) $\tilde{X}_i$ from the x_i (in order to generate branches of the UST according to Wilson's algorithm), at least one of them will hit $\partial B(N)$ before γ. In other words, we get the following bound :

$$\mathbb{P}(\mathcal{A}(L, N)) \leq \mathbb{E}(K) \times \mathbb{E}\left(\sup_{x \in \partial B(3L)} \mathbb{P}_x(\tilde{X} \cap \gamma = \emptyset | \mathcal{F}) \right).$$

Let us now decompose the path of $\tilde{X}$ according to its down- and up-crossings of the annulus between $\partial B(4L)$ and $\partial B(3L)$. At each down-crossing, it has a probability bounded from below (say by a constant b) to disconnect $B(3L)$ from $\partial B(4L)$, and therefore to hit γ. We can therefore decompose the path X_i according to its number N of downcrossings, and see immediately that

$$\sup_{x \in \partial B(3L)} \mathbb{P}_x(\tilde{X} \cap \gamma = \emptyset | \mathcal{F}) \leq \frac{1}{1-b} \sup_{x \in \partial B(4L)} \mathbb{P}_x(X \cap \gamma = \emptyset \text{ and } N = 0 | \mathcal{F}),$$

where X is a simple random walk stopped at its first hitting of $\partial B(N)$. Let us call γ^{5L} the part of γ after its last hitting time of $\partial B(5L)$. We trivially have, for an arbitrary point $x \in \partial B(4L)$ and any given path γ, that $\mathbb{P}_x(X \cap \gamma = \emptyset \text{ and } N = 0) \leq \mathbb{P}_x(X \cap \gamma^{5L} = \emptyset)$. Using (1), we get the uniform bound

$$\mathbb{P}(\mathcal{A}(L, N)) \leq \frac{\mathbb{E}(K)}{1-b} \mathbb{E}\left(\sup_{x \in \partial B(4L)} \mathbb{P}_x(X \cap \gamma^{5L} = \emptyset | \mathcal{F}) \right) \leq C'(5L/N)^{3/4}$$

where C' is a universal constant, which wraps up the proof. ☐

4.3. **Arm-estimates imply Lemma 8 and Lemma 7.**

Proof of Lemma 8. We can bound the expected value of the renormalized length L^δ by $\delta^{5/4}$ times the sum over all pairs of edges e_0 (in the dual lattice) and e_1 (in the original lattice) that are at distance at most ε of each other of the probability of the following intersection of events $E(e_0, e_1)$:

- The edge e_0 belongs to the dual cycle $\mathcal{C}^\delta$ (that appears when closing the edge e).
- The edge e_0 is at distance greater than η from e, and in the ball of radius $1/\eta$ around the origin.
- If we erase the two edges e and e_1 from the UST, the edge e_0 is no longer on the interface between the clusters that contain z_0 and z_1.
- The edge e_1 is cut out during the cutting procedure (note that this event occurs independently of the rest, with probability $\delta^{5/4}$ times a constant that depends on t).

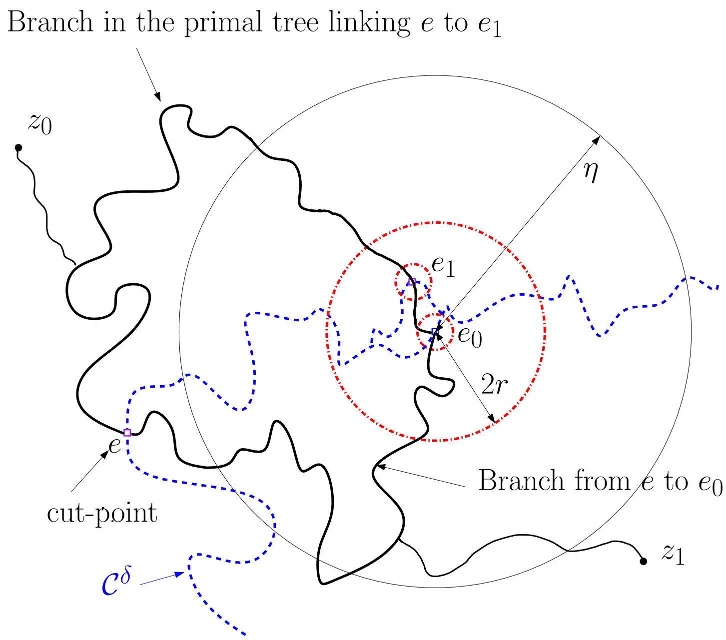

FIGURE 10. Arm events appearing in $E(e_0, e_1)$

For any f edge in $\delta\mathbb{Z}^2$ and $l_1 \leq l_2$, denote by $\mathcal{A}_f(l_1, l_2)$ the four arms event in the annulus $B(x, l_2) \setminus B(x, l_1)$ centered at the middle point x of the edge f for the UST on $\delta\mathbb{Z}^2$. We have that, if we set $r := d(e_0, e_1)$, then (see Fig. 10):

$$E(e_0, e_1) \subset \mathcal{A}_{e_0}(\delta/2, r/3) \cap \mathcal{A}_{e_1}(\delta/2, r/3) \cap \mathcal{A}_{e_0}(2r, \eta).$$

Using the upper bounds on the probabilities of these events given by Lemma 9, together with the fact that the bounds on the first two are independent of the boundary conditions (so it is possible to first condition on the last one, and then to bound the conditional probability of the first two), we get readily that

$$\begin{aligned}\mathbb{E}(L^\delta) &\leq \delta^{5/4} \times C(t)\delta^{5/4} \times \sum_{e_0 \in B(0,1/\eta)} \sum_{e_1 \in B(e_0,\varepsilon)} (r/\delta)^{-3/4} \times (r/\delta)^{-3/4} \times (\eta/r)^{-3/4} \\ &\leq C(\eta)\delta^{5/4+5/4+3/4+3/4-2} \times \sum_{x \in \delta\mathbb{Z}^2 \cap B(0,\varepsilon)\setminus\{0\}} |x|^{-3/4} \leq C(\eta)\varepsilon^{5/4}.\end{aligned}$$

□

Proof of Lemma 7. The proof goes along similar lines than the previous one. We can bound the expected renormalized length by $\delta^{5/4}$ times the sum over all pairs of edges e_0 (in the original lattice) and e_1 (in the dual lattice) such that $r := d(e_0, e_1) \leq \eta$ of the probability of the intersection of the following events:

- The edge e_0 is η_0-inside the UST branch from z_0 to z_1, and it is removed by the cutting procedure.
- The edge e_1 belongs to the dual cycle that appears when closing the edge e_0.

As in the previous argument, we can note that this event (for given e_0 and e_1) is included in the joint occurrence of four arms events $\mathcal{A}_{e_0}(\delta/2, r/3) \cap \mathcal{A}_{e_1}(\delta/2, r/3) \cap \mathcal{A}_{e_0}(2r, \eta_0)$, and we can conclude using the same computation. □

Appendix A. Strong convergence of Uniform Spanning Trees

In this appendix, we show that USTs with different boundary conditions (wired, free, whole plane) converge in the scaling limit, in the sense that their finite subtrees parametrized by length converge. This is done here by pushing the result by Lawler and Viklund about convergence of chordal LERW towards SLE_2 in its natural parametrization (this result from [19] is stated in this appendix as Result F) to the other setups by (local) absolute continuity between the chordal and other settings, using loop-soups to evaluate the Radon-Nikodym derivatives and the fact that discrete Green's functions approximate well their continuous counterparts in the scaling limit (this in the spirit of the more probabilistic presentation in [15] for instance).

A.1. **Random walk and loop measure.** Recall from paragraph 2.2 that Ω is either the entire plane or a bounded simply connected domain with analytic boundary and Ω^δ is a simply connected discretization of it at mesh size δ defined as follows: Fix $\xi \in \Omega$ and consider A the simply connected subset of $\delta\mathbb{Z}^2 \cap \Omega$ containing ξ. The set Ω^δ is the union of the squares $\{a + x + \mathbf{i}\, y,\ |x| \leq 1/2,\ |y| \leq 1/2\}$ over $a \in A$.

In this appendix, we denote by $\mu^{\mathrm{RW}}_{\Omega^\delta}$ is the measure on RW path in the domain Ω^δ, i.e. the measure that assigns to a finite path X the probability $(1/4)^{\sharp X}$, where $\sharp X$ is the number of steps of the path. Let $\mu^{e}_{\Omega^\delta}$ denote the (oriented) excursion measure, i.e. the measure on RW paths in Ω^δ restricted to those paths having their two extremities on the boundary of the domain and otherwise staying in the interior.

Moreover, let $\mu^{\ell}_{\Omega^\delta}$ and $\mu^{r\ell}_{\Omega^\delta}$ respectively denote the (oriented, unrooted) RW loop measure in Ω^δ, and the reflected RW loop measure in Ω^δ. In these measures, a loop X comes with a weight $(1/4)^{\sharp X}$, together with a factor $1/J$ if the loop is the concatenation of J times the same loop (see [26]).

A.2. **Radial and full-plane LERWs.** Let us first state the result from Lawler and Viklund that we are going to use as the main building block:

Result F ([19])**.** Let Ω be a bounded simply connected domain with analytic boundary, and two marked boundary points a, b. Let $a^\delta, b^\delta \in \delta\mathbb{Z}^2$ be approximations of a, b in $\mathbb{C}\backslash\Omega^\delta$ such that there are edges $[a^\delta, a'], [b^\delta, b']$ of $\delta\mathbb{Z}^2$ with $a', b' \in \Omega^\delta$.

The chordal LERW from a^δ to b^δ in Ω^δ converges to chordal SLE_2 when both are naturally parametrized.

We want to deduce the following Proposition (otherwise refered to as Result C in the paper):

Proposition 10. *Let Ω be a bounded simply connected domain with analytic boundary, let a be a marked boundary point, x_0 a marked interior point. The domain is approximated by lattice domains $(\Omega^\delta, a^\delta, x_0^\delta)$ as above. Denote by $\gamma^\delta : [0, s^\delta] \to \Omega^\delta$ the radial LERW from a^δ to x_0^δ in Ω^δ and $\gamma : [0, s] \to \Omega$ the radial SLE_2 from a to x_0 in its natural parametrization. We extend the curves until $s \vee s^\delta$ by keeping them at their ending point so that they are defined on the same time-interval. Then, the curve $t \mapsto \gamma(t/\delta^{5/4})$ converges in distribution to γ for the topology of the supremum norm.*

Proof. In the discretization Ω^δ of a bounded simply-connected domain Ω, we consider a chordal LERW between two boundary points a^δ and b^δ, and a radial LERW from a^δ to an interior marked point x_0^δ.

Given a simple discrete path $\gamma^\delta : [0, \tau] \to \Omega^\delta$ from a^δ to an interior point $y^\delta = \gamma^\delta(\tau)$ which does not hit x_0^δ (approximating a continuous simple path γ), the probability that the radial LERW starts by γ^δ is given by the sum over all paths X in $\Omega^\delta\backslash\{x_0\}$ whose loop erasure $LE(X)$ equals γ^δ of the probability that a simple random walk (in the plane) starting at a^δ, conditioned to enter immediately in Ω^δ and hit x_0^δ before $\partial\Omega^\delta$, starts by X and then hits x_0^δ before $\partial(\Omega^\delta \setminus \gamma^\delta)$. Therefore,

if $\sum_{X:a^\delta\to y^\delta}$ denotes the sum over all nearest neighbor paths that start at a^δ and end at y^δ, and if $x_0^\delta \notin X$ denotes the event that X does not hit the point x_0^δ,

$$\begin{aligned}
\mathbb{P}^r[\gamma^\delta] &= \frac{\sum\limits_{X:a^\delta\to y^\delta} 4^{-\sharp X}\, \mathbb{1}_{x_0^\delta\notin X}\, \mathbb{1}_{LE(X)=\gamma^\delta}\, \mathbb{P}[\text{RW from } y^\delta \text{ touches } x_0^\delta \text{ before } \partial(\Omega^\delta\setminus\gamma^\delta)]}{\mathbb{P}[\text{RW from } a^\delta \text{ touches } x_0^\delta \text{ before } \partial\Omega^\delta]} \\
&= \frac{\sum\limits_{X:a^\delta\to y^\delta} 4^{-\sharp X}\, \mathbb{1}_{x_0^\delta\notin X}\mathbb{1}_{LE(X)=\gamma^\delta}\, \mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e: y^\delta\to x_0^\delta\})}{\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e: a^\delta\to x_0^\delta\})}.
\end{aligned}$$

Similarly, the probability that the chordal LERW starts by γ^δ reads:

$$\begin{aligned}
\mathbb{P}^c[\gamma^\delta] &= \frac{\sum\limits_{X:a^\delta\to y^\delta} 4^{-\sharp X}\mathbb{1}_{LE(X)=\gamma^\delta}\, \mathbb{1}_{x_0^\delta\notin X}\, \mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e: y^\delta\to b^\delta\})}{\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e: a^\delta\to b^\delta\})} \\
&\quad + \frac{\sum\limits_{X:a^\delta\to y^\delta} 4^{-\sharp X}\mathbb{1}_{LE(X)=\gamma^\delta}\, \mathbb{1}_{x_0^\delta\in X}\, \mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e: y^\delta\to b^\delta\})}{\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e: a^\delta\to b^\delta\})}.
\end{aligned}$$

The second term in the RHS is negligible. Indeed, its relative contribution to the RHS can be expressed in term of the loop measure $\mu^\ell_{\Omega^\delta}$ (see [26], Proposition 1.3):

$$\frac{\sum\limits_{X:a^\delta\to y^\delta} 4^{-\sharp X}\mathbb{1}_{LE(X)=\gamma^\delta}\mathbb{1}_{x_0^\delta\in X}}{\sum\limits_{X:a^\delta\to y^\delta} 4^{-\sharp X}\mathbb{1}_{LE(X)=\gamma^\delta}} = \frac{4^{-|\gamma^\delta|}\exp\left(\mu^\ell_{\Omega^\delta}(\{\ell, \ell\cap\gamma^\delta\neq\emptyset\ ,\ x_0^\delta\in\ell\})\right)}{4^{-|\gamma^\delta|}\exp\left(\mu^\ell_{\Omega^\delta}(\{\ell, \ell\cap\gamma^\delta\neq\emptyset\})\right)},$$

and the mass of the loops passing through x_0^δ divided by the total mass of loops tends to zero when $\delta\to 0$. Hence we can express the Radon-Nikodym derivative of radial LERW with respect to chordal LERW as:

$$\frac{\mathbb{P}^r[\gamma^\delta]}{\mathbb{P}^c[\gamma^\delta]} = \frac{\mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e: \gamma^\delta(\tau)\to x_0^\delta\})\ \mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e: a^\delta\to b^\delta\})}{\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e: a^\delta\to x_0^\delta\})\ \mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e: \gamma^\delta(\tau)\to b^\delta\})} \times (1+o(1)). \tag{2}$$

We now rewrite this ratio in terms of harmonic quantities. For this, let us introduce a cut c that disconnects a from both b and x_0 in the domain Ω. We also consider c^δ a discrete approximation of this cut. We call Ω_c (resp. Ω_c^δ) the subdomain of $\Omega\setminus c$ (resp. $\Omega^\delta\setminus c^\delta$) containing a (resp. a^δ). Similarly, let us introduce a cut d that disconnects $\gamma(\tau)$ from both b and x_0 in the domain $\Omega\setminus\gamma$. We choose a discrete approximation d^δ of the cut d. We call $(\Omega\setminus\gamma)_d$ (resp. $(\Omega^\delta\setminus\gamma^\delta)_d$) the subdomain of $\Omega\setminus(\gamma\cup d)$ (resp. $\Omega^\delta\setminus(\gamma^\delta\cup d)$) containing $\gamma(\tau)$ (resp. $\gamma^\delta(\tau)$) (see Figure 11).

We also pick, for normalization purposes, two subintervals I_c and I_d on the boundaries $\partial\Omega_c\setminus c$ and $\partial(\Omega\setminus\gamma)_d\setminus d$ respectively, as well as a point $x_1\in\Omega\setminus(\gamma\cup\{x_0\})$. We also pick discrete approximations I_c^δ and I_d^δ and x_1^δ of these three objects.

Let $L_{a^\delta}^{\Omega_c^\delta}$ be the measure on endpoints of RW excursions in Ω_c^δ, starting at a^δ (with the natural discrete normalization, i.e. such that the total measure of the boundary is one). We also call

$$L_{a^\delta}^{\Omega_c^\delta, I_c^\delta} = \frac{L_{a^\delta}^{\Omega_c^\delta}(\cdot)}{L_{a^\delta}^{\Omega_c^\delta}(I_c^\delta)}$$

the endpoint measure normalized so that $L_{a^\delta}^{\Omega_c^\delta, I_c^\delta}(I_c^\delta)=1$.

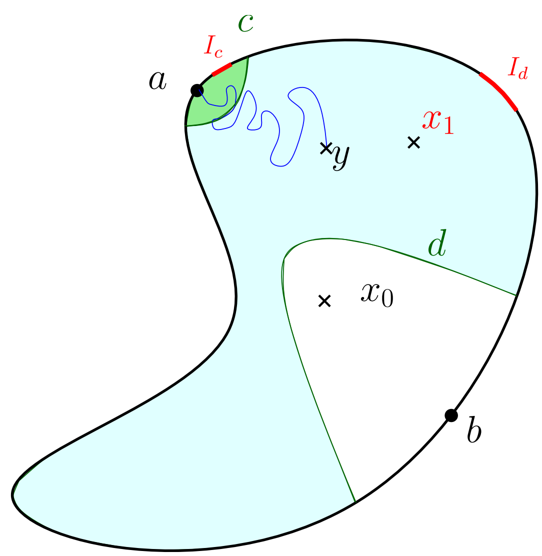

FIGURE 11. Representation of Ω, γ and the other quantities

Let $m_x^{\Omega^\delta}$ be the harmonic measure on the boundary of Ω^δ seen from an interior point x (of total mass one). The quantity

$$P_{b^\delta}^{\Omega^\delta, x_1^\delta}(\cdot) = \frac{m_\cdot^{\Omega^\delta}(\{b^\delta\})}{m_{x_1^\delta}^{\Omega^\delta}(\{b^\delta\})}$$

is then the Poisson kernel in Ω^δ with singularity at b^δ, normalized at the interior point x_1^δ.

Let $G^{\Omega^\delta}(\cdot, x_0^\delta)$ be the Green's function in Ω^δ normalized such that $G^{\Omega^\delta}(x_0^\delta, x_0^\delta) = 1$. We also let

$$G_{x_1^\delta}^{\Omega^\delta}(\cdot, x_0^\delta) = \frac{G^{\Omega^\delta}(\cdot, x_0^\delta)}{G^{\Omega^\delta}(x_1^\delta, x_0^\delta)}$$

be the Green's function normalized such that $G_{x_1^\delta}^{\Omega^\delta}(x_1^\delta, x_0^\delta) = 1$.

By stopping RW paths from a^δ until their first visit of the cut c^δ (or d^δ), we see that we can rewrite:

$$\begin{aligned}
\mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e : y^\delta \to x_0^\delta\}) &= \int_{x\in d^\delta} G^{\Omega^\delta\setminus\gamma^\delta}(x, x_0^\delta) L_{y^\delta}^{(\Omega^\delta\setminus\gamma^\delta)_d}(dx) \\
\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e : a^\delta \to x_0^\delta\}) &= \int_{x\in c^\delta} G^{\Omega^\delta}(x, x_0^\delta) L_{a^\delta}^{\Omega_c^\delta}(dx) \\
\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e : a^\delta \to b^\delta\}) &= \int_{x\in c^\delta} m_x^{\Omega^\delta}(\{b^\delta\}) L_{a^\delta}^{\Omega_c^\delta}(dx) \quad \text{and} \\
\mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e : y^\delta \to b^\delta\}) &= \int_{x\in d^\delta} m_x^{\Omega^\delta\setminus\gamma^\delta}(\{b^\delta\}) L_{y^\delta}^{(\Omega^\delta\setminus\gamma^\delta)_d}(dx).
\end{aligned}$$

We need to use renormalized versions of the harmonic quantities L, m and G in order to get convergence in the scaling limit. The ratio (2) will make each of our four normalization choices cancel out.

Indeed, let us write down:

$$\begin{aligned}
\mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e: y^\delta \to x_0^\delta\}) &= L_{y^\delta}^{(\Omega^\delta\setminus\gamma^\delta)_d}(I_d^\delta)\; G^{\Omega^\delta\setminus\gamma^\delta}(x_1^\delta, x_0^\delta)\int_{x\in d^\delta} G_{x_1^\delta}^{\Omega^\delta\setminus\gamma^\delta}(x, x_0^\delta) L_{y^\delta}^{(\Omega^\delta\setminus\gamma^\delta)_d, I_d^\delta}(dx)\\
\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e: a^\delta \to x_0^\delta\}) &= L_{a^\delta}^{\Omega_c^\delta}(I_c^\delta)\; G^{\Omega^\delta}(x_1^\delta, x_0^\delta)\int_{x\in c^\delta} G_{x_1^\delta}^{\Omega^\delta}(x, x_0^\delta) L_{a^\delta}^{\Omega_c^\delta, I_c^\delta}(dx)\\
\mu^{\mathrm{RW}}_{\Omega^\delta}(\{e, e: a^\delta \to b^\delta\}) &= L_{a^\delta}^{\Omega_c^\delta}(I_c^\delta)\; m_{x_1^\delta}^{\Omega^\delta}(\{b^\delta\})\int_{x\in c^\delta} P_{b^\delta}^{\Omega^\delta, x_1^\delta}(x) L_{a^\delta}^{\Omega_c^\delta, I_c^\delta}(dx) \ \text{ and}\\
\mu^{\mathrm{RW}}_{\Omega^\delta\setminus\gamma^\delta}(\{e, e: y^\delta \to b^\delta\}) &= L_{\gamma^\delta(\tau)}^{(\Omega^\delta\setminus\gamma^\delta)_d}(I_d^\delta)\; m_{x_1^\delta}^{\Omega^\delta\setminus\gamma^\delta}(\{b^\delta\})\int_{x\in d^\delta} P_{b^\delta}^{\Omega^\delta\setminus\gamma^\delta, x_1^\delta}(x) L_{y^\delta}^{(\Omega^\delta\setminus\gamma^\delta)_d, I_d^\delta}(dx).
\end{aligned}$$

We can then rewrite (2) as:

$$(3)\quad \begin{aligned}
\frac{\mathbb{P}^r[\gamma^\delta]}{\mathbb{P}^c[\gamma^\delta]} = {} & \frac{G^{\Omega^\delta\setminus\gamma^\delta}(x_1^\delta, x_0^\delta)}{G^{\Omega^\delta}(x_1^\delta, x_0^\delta)} \times \frac{m_{x_1^\delta}^{\Omega^\delta}(\{b^\delta\})}{m_{x_1^\delta}^{\Omega^\delta\setminus\gamma^\delta}(\{b^\delta\})} \times \\
& \frac{\int_{x\in d^\delta} G_{x_1^\delta}^{\Omega^\delta\setminus\gamma^\delta}(x, x_0^\delta) L_{y^\delta}^{(\Omega^\delta\setminus\gamma^\delta)_d, I_d^\delta}(dx) \int_{x\in c^\delta} P_{b^\delta}^{\Omega^\delta, x_1^\delta}(x) L_{a^\delta}^{\Omega_c^\delta, I_c^\delta}(dx)}{\int_{x\in c^\delta} G_{x_1^\delta}^{\Omega^\delta}(x, x_0^\delta) L_{a^\delta}^{\Omega_c^\delta, I_c^\delta}(dx) \int_{x\in d^\delta} P_{b^\delta}^{\Omega^\delta\setminus\gamma^\delta, x_1^\delta}(x) L_{y^\delta}^{(\Omega^\delta\setminus\gamma^\delta)_d, I_d^\delta}(dx)} \times (1 + o(1)).
\end{aligned}$$

All four integrals in (3) are expressed in terms of discrete harmonic quantities that are normalized to naturally converge to their continuous counterparts (see e.g. [7]).

Let us now show that the ratio

$$(4)\qquad \frac{G^{\Omega^\delta\setminus\gamma^\delta}(x_1^\delta, x_0^\delta)}{G^{\Omega^\delta}(x_1^\delta, x_0^\delta)}.$$

converges. We want to show that the renormalization factor that makes the function $x \mapsto G^{\Omega^\delta}(x, x_0^\delta)$ converge can be chosen independently of the domain Ω^δ, so that we can pick the same factor to make the function $x \mapsto G^{\Omega^\delta\setminus\gamma^\delta}(x, x_0^\delta)$ converge. The discrete Green function converges in the classical normalization G_c (such that $\Delta_x G_c^{\Omega^\delta}(x, x_0^\delta)_{|x=x_0^\delta} = 1$) towards a function with a logarithmic singularity at x_0, which does not depend on the domain Ω. The ratio $r_c(x) = \frac{G_c^{\Omega\setminus\gamma}(x,x_0)}{G_c^{\Omega}(x,x_0)}$ can thus be extended to $x = x_0$ by $r_c(x_0) = 1$, and is hence equal to the limit of the ratio $r^\delta(x)$ (normalized in the discrete so that $r^\delta(x_0) = 1$)

$$r^\delta(x) = \frac{G^{\Omega^\delta\setminus\gamma^\delta}(x, x_0^\delta)}{G^{\Omega^\delta}(x, x_0^\delta)}.$$

In particular, the ratio (4), being equal to the value $r^\delta(x_1^\delta)$, converges to $r_c(x_1)$.

The convergence of the ratio

$$(5)\qquad \frac{m_{x_1^\delta}^{\Omega^\delta}(\{b^\delta\})}{m_{x_1^\delta}^{\Omega^\delta\setminus\gamma^\delta}(\{b^\delta\})}$$

can be obtained in the following way. Let us consider a secondary cut $\tilde{d}$ (with discrete approximation $\tilde{d}^\delta$) that disconnects d from the boundary point b. We call D (and similarly D^δ) the connected component of $\Omega \setminus d$ with the point b on its boundary. In contrast, we call $\widetilde{D}$ (and similarly $\widetilde{D}^\delta$) the connected component of $\Omega \setminus \tilde{d}$ which does not have the point b on its boundary. Finally, we pick a point $y_0 \in \tilde{d}$, and approximations y_0^δ of it.

We now show that the ratios

$$\frac{m_{x_1^\delta}^{\Omega^\delta}(\{b^\delta\})}{m_{y_0^\delta}^{D^\delta}(\{b^\delta\})} \text{ and } \frac{m_{x_1^\delta}^{\Omega^\delta\setminus\gamma^\delta}(\{b^\delta\})}{m_{y_0^\delta}^{D^\delta}(\{b^\delta\})}$$

converge. We only focus on the first quantity as the second follows similarly.

Given a point $y \in \tilde{d}^\delta$, let us call $T^\delta(y,\cdot)$ the measure on $\tilde{d}^\delta$ given by

$$T^\delta(y,\cdot) = \sum_{z\in d^\delta} m_z^{\tilde{D}^\delta}(\cdot) m_y^{D^\delta}(\{z\}).$$

This measure-valued function converges to its continuous equivalent defined thus:

$$T(y,\cdot) = \int_{z\in d} m_z^{\tilde{D}^\delta}(\cdot) m_y^{D^\delta}(dz).$$

We can now express the harmonic measure $m_{x_1^\delta}^{\Omega^\delta}(\{b^\delta\})$ by decomposing the path of a random walk from x_1^δ to b^δ into its n trips back and forth between the two cuts d and $\tilde{d}$:

$$\frac{m_{x_1^\delta}^{\Omega^\delta}(\{b^\delta\})}{m_{y_0^\delta}^{D^\delta}(\{b^\delta\})} = \sum_{n\geq 1}\sum_{y_1,\cdots,y_n\in\tilde{d}} \frac{m_{y_n}^{D^\delta}(\{b^\delta\})}{m_{y_0^\delta}^{D^\delta}(\{b^\delta\})}\left(\prod_{i=1}^{n-1} T^\delta(y_i, dy_{i+1})\right) m_{x_1^\delta}^{\tilde{D}^\delta}(dy_1).$$

Recall that $\frac{m_{y_n}^{D^\delta}(\{b^\delta\})}{m_{y_0^\delta}^{D^\delta}(\{b^\delta\})} = P_{b^\delta}^{D^\delta, y_0^\delta}(y_n)$, and hence all terms in the above expression converge to their continuous counterparts. As explained above, this yields the convergence of the ratio (5).

We have now shown that all terms in the Radon-Nikodym derivative expression (3) converge; their product therefore converges as well.

Now, note that the natural parametrization are functions of the curves, and that we know that radial LERW converges to radial SLE_2 in the capacity parametrization. The claim therefore follows by Result F and convergence of the Radon-Nikodym derivatives of radial LERW with respect to chordal LERW. $\square$

Corollary 11 (Wired UST). *Let Ω be a bounded simply connected domain with analytic boundary, let $z_1, \cdots, z_n$ be n points in Ω at a positive distance from $\partial\Omega$ and for each $\delta > 0$ let $z_1^\delta, \cdots, z_n^\delta$ be approximations of these points on $\delta\mathbb{Z}^2$. Consider the branches in the* ***wired UST*** *γ_1^δ linking $\partial\Omega^\delta$ to z_1^δ and γ_i^δ : $\partial\Omega^\delta \cup \gamma_1^\delta \cup \cdots \cup \gamma_{i-1}^\delta$ to z_i^δ for $i \in \{2, \cdots n\}$ and their continuous counterparts γ_i and adapt their length as in Proposition 10. Then the branches $t \mapsto \gamma_i^\delta(t/\delta^{5/4})$ jointly converges to γ_i for the topology of the supremum norm when $\delta \to 0$.*

As a consequence, the tree $\mathcal{T}_{z_1,\ldots,z_n,\partial\Omega}^\delta$ (defined by the union of the branches γ_i^δ, $i \in \{1, \cdots, n\}$) parametrized by its arc-length from its boundary and renormalized by $\delta^{5/4}$ converges to its continuous counterpart $\mathcal{T}_{z_1,\ldots,z_n,\partial\Omega}$.

Proposition 12 (Planar UST). *Let $z_1, \cdots z_n$ be $n \geq 3$ points in the plane and $z_1^\delta, \cdots z_n^\delta$ be approximations of those points in $\delta\mathbb{Z}^2$. Define by γ_1^δ the branch between z_1^δ and z_2^δ in the* ***plane UST*** *of $\delta\mathbb{Z}^2$ and for all $i \in \{2, \cdots, n\}$, by γ_i^δ the branch between z_i^δ and $\gamma_1^\delta \cup \cdots \cup \gamma_{i-1}^\delta$. Let z_0^δ be the intersection point between the branches γ_1^δ and γ_2^δ. The tree $\mathcal{T}_{z_1,\ldots,z_n}^\delta$ defined by the union of the n branches γ_i^δ parametrized by its arc-length from z_0^δ and renormalized by $\delta^{5/4}$ converges towards its continuous counterpart.*

Proof. Let us fix a large scale $R > 0$. The law of the full-plane LERW (from 0) until its first exit time of $B(0,R)$ (the ball of radius R) is absolutely continuous with respect to the law of radial LERW (from 0) in the ball of radius $4R$ until its first exit time of the ball of radius R ([21, Corollary

4.5]). Hence, the time-reversed full plane LERW from $\partial B(0,R) \to 0$ in $B(0,R)$ (parametrized by its length divided by $\gamma^{5/4}$) converges (for the topology of the sup-norm) towards its continuous counterpart. As we only need convergence of length of branches in finite subtrees, and away from their extremities, this completes the proof. $\square$

A.3. **The free UST.**

Proposition 13 (Free UST). *Let Ω be a bounded simply connected domain with analytic boundary and Ω^δ its approximation. Let $z_1, \cdots z_n$ be $n \geq 3$ points in Ω at a positive distance from the boundary and $z_1^\delta, \cdots z_n^\delta$ be approximation of those points in $\delta\mathbb{Z}^2$. Define by γ_1^δ the branch in the* ***free UST*** *between z_1^δ and z_2^δ and for all $i \in \{2, \cdots, n\}$, by γ_i^δ the branch between z_i^δ and $\gamma_1^\delta \cup \cdots \cup \gamma_{i-1}^\delta$. Let z_0^δ be the intersection point between the branches γ_1^δ and γ_2^δ. The tree $\mathcal{T}^\delta_{z_1,\ldots,z_n}$ defined by the union of the n branches γ_i^δ parametrized by its arc-length from z_0^δ renormalized by $\delta^{5/4}$ converges towards its continuous counterpart.*

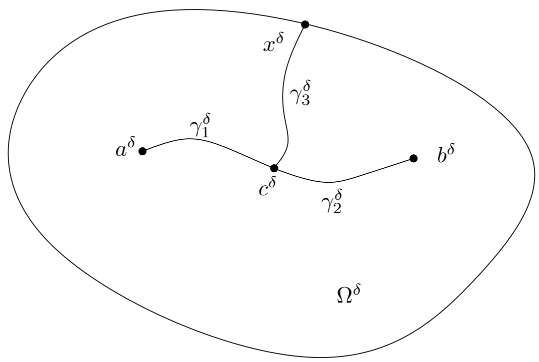

FIGURE 12. The wired branch $\gamma^\delta = \gamma_1^\delta \cup \gamma_2^\delta$ from a^δ to b^δ is produced via Wilson's algorithm. This algorithm also generates the branch γ_3^δ from c^δ to $x^\delta \in \partial\Omega^\delta$ linking the branch from a^δ to b^δ to the wired boundary.

Proof. By induction, it suffices to compare the probability to have a given path between two points a and b in the free and wired UST in Ω. Recall that, when sampling a UST via Wilson's algorithm, the erased loops have the law of a Poisson process of oriented Random Walk loops.

Hence, the probability to see a path γ^δ as the branch from a^δ to b^δ in the free UST in Ω^δ is given by (see [26], Proposition 1.3):

$$\begin{aligned}\mathbb{P}^f[\gamma^\delta] &= \frac{1}{Z} 4^{-|\gamma^\delta|} \exp\left(\mu^{r\ell}_{\Omega^\delta}(\{\ell, \ell \cap (\gamma^\delta \setminus \{b^\delta\}) \neq \emptyset\})\right) \\ &= \frac{1}{Z'} 4^{-|\gamma^\delta|} \exp\left(\mu^{r\ell}_{\Omega^\delta}(\{\ell, \ell \cap (\gamma^\delta \setminus \{b^\delta\}) \neq \emptyset \text{ and } \ell \text{ does not disconnect } a^\delta \text{ from } b^\delta\})\right).\end{aligned}$$

And the probability $\mathbb{P}^w[\gamma^\delta]$ to see a path γ^δ as the branch from a^δ to b^δ in the wired UST in Ω^δ is given by (see Figure 12):

$$\frac{1}{Z_w} \sum_{\substack{c^\delta \in \gamma^\delta \\ \gamma_3^\delta : c^\delta \to \partial\Omega^\delta}} 4^{-|\gamma_1^\delta| - |\gamma_3^\delta| - |\gamma_2^\delta|} \exp\left(\mu^{\ell}_{\Omega^\delta}(\{\ell, \ell \cap (\gamma_1^\delta \cup \gamma_3^\delta) \neq \emptyset\})\right) \exp\left(\mu^{\ell}_{\Omega^\delta \setminus (\gamma_1^\delta \cup \gamma_3^\delta)}(\{\ell, \ell \cap (\gamma_2^\delta) \neq \emptyset\})\right)$$

which can be rewritten as

$$\begin{aligned}
\mathbb{P}^w[\gamma^\delta] &= \frac{1}{Z_w} 4^{-|\gamma^\delta|} \exp\left(\mu^\ell_{\Omega^\delta}(\{\ell, \ell \cap \gamma^\delta \neq \emptyset\})\right) \sum_{c^\delta, \gamma_3^\delta} 4^{-|\gamma_3^\delta|} \exp\left(\mu^\ell_{\Omega^\delta \setminus \gamma^\delta}(\{\ell, \ell \cap (\gamma_3^\delta) \neq \emptyset\})\right) \\
&= \frac{1}{Z_w} 4^{-|\gamma^\delta|} \exp\left(\mu^\ell_{\Omega^\delta}(\{\ell, \ell \cap \gamma^\delta \neq \emptyset\})\right) \sum_{c^\delta \in \gamma^\delta, x^\delta \in \partial\Omega^\delta} \mu^e_{\Omega^\delta \setminus \gamma^\delta}(\{e, e \text{ from } x^\delta \text{ to } c^\delta\}) \\
&= \frac{1}{Z_w} 4^{-|\gamma^\delta|} \exp\left(\mu^\ell_{\Omega^\delta}(\{\ell, \ell \cap \gamma^\delta \neq \emptyset\})\right) \mu^e_{\Omega^\delta \setminus \gamma}(\{e, e \text{ from } \partial\Omega^\delta \text{ to } \gamma^\delta\}) \\
&= \frac{1}{Z_w} 4^{-|\gamma^\delta|} \exp\left(\mu^\ell_{\Omega^\delta}(\{\ell, \ell \cap \gamma^\delta \neq \emptyset\})\right) \frac{1}{2} \mu^e_{\Omega^\delta}(\{e, e \cap \gamma^\delta \neq \emptyset\}).
\end{aligned}$$

Hence,

$$\begin{aligned}
\log \frac{\mathbb{P}^f[\gamma^\delta]}{\mathbb{P}^w[\gamma^\delta]} &= C + \mu^{r\ell}_{\Omega^\delta}(\{\ell, \ell \cap (\gamma^\delta \setminus \{b^\delta\}) \neq \emptyset \text{ and } \ell \text{ does not disconnect } a^\delta \text{ from } b^\delta\}) \\
&- \mu^\ell_{\Omega^\delta}(\{\ell, \ell \cap \gamma^\delta \neq \emptyset\}) - \log\left(\mu^e_{\Omega^\delta}(\{e, e \cap \gamma^\delta \neq \emptyset\})\right),
\end{aligned}$$

where C is a constant, $\mu^{r\ell}_{\Omega^\delta} - \mu^\ell_{\Omega^\delta}$ allows for the cancellation of microscopic loops touching γ^δ. Let us note here that the mass $\mu^{r\ell}_{\Omega^\delta}(\{\ell, b^\delta \in \ell \text{ and } \ell \text{ does not disconnect } a^\delta \text{ from } b^\delta\})$ tends to zero in the scaling limit.

As the discrete measures $\mu^\ell_{\Omega^\delta}$, $\mu^{r\ell}_{\Omega^\delta}$ and $\mu^e_{\Omega^\delta}$ converge towards equivalent measures for Brownian motion (Lemmas 14, 15 and 16), the claim follows from the result for the wired UST (Corollary 11) by absolute continuity. □

Lemma 14. *The measure $\mu^\ell_{\Omega^\delta}$ converges to its continuous counterpart.*

Proof. We refer to [18], which contains a much finer convergence result. □

Lemma 15. *The measure $\mu^e_{\Omega^\delta}$ converges to its continuous counterpart.*

Proof. The discrete excursion measure can be constructed from the measures $\mu^{\mathrm{RW}}_{z\to\partial\Omega^\delta}$ on random walks from an interior point z to the boundary of the domain in the following way. For each interior point z, take two independent copies random walks sampled according to $\mu^{\mathrm{RW}}_{z\to\partial\Omega^\delta}$, and concatenate the reversed first path to the second one in order to get a path that goes from the boundary $\partial\Omega^\delta$ to itself via z. Divide the mass by the total time $n = T\delta^{-2}$ (which is the length of the discrete path). The excursion measure is then obtained by summing over z.

The same relationship between RW measure and excursion measure holds in the continuous setting: if one reweights the excursion measure by the total time T, and picks a time t uniformly at random on it (corresponding to a point z on its path), then by the Markov property, we see that the future is a BM from z to $\partial\Omega$ independent of the past, conditionally on z.

The convergence of excursion measures $\mu^e_{\Omega^\delta}$ follows from the convergence of the random walk measures. □

Lemma 16. *For all polygonal domain Ω (whose boundaries are formed of edges of the lattice $\alpha\mathbb{Z}^2$ of mesh size α), the measure $\mu^{r\ell}_{\Omega^\delta}$ converges to its continuous counterpart. Therefore, for all domain Ω, the measure $\mu^{r\ell}_{\Omega^\delta}$ restricted to loops visiting a (fixed) compact in the interior of Ω converges as well.*

Proof. Let Ω be a polygonal domain of mesh size α, and consider discrete approximation Ω^δ of Ω at mesh sizes δ dividing α. We can then do exact reflections of the domain Ω near boundary segments and corners.

More precisely, let us fix scales ϵ and ϵ'. We call ϵ, ϵ'-admissible loops, the loops that are macroscopic (diameter larger than ϵ), not dense (e.g. do not disconnect two marked interior points a and b) and that visit an ϵ'-neighborhood of one of the corners (we actually only to exclude concave corners). The mass of ϵ, ϵ'-admissible loops tends to 0 when ϵ' goes to 0 (for a fixed ϵ).

We now consider ϵ, ϵ'-admissible loops rooted at a point z. We introduce a sequence of stopping times τ_n, σ_n on our loop: call σ_0 the starting time, then time τ_n is the first hitting time of the $\epsilon'/2$-neighborhood of a boundary segment after σ_{n-1}, and σ_n is the first exit time of the ϵ'-neighborhood of the boundary segment corresponding to τ_n. We call N the index of the last time σ before the loop ends. All subpaths of the loop between consecutive stopping times converge nicely. Indeed, note that if $N = 0$, then Lemma 14 allows us to conclude. Otherwise, note that all but the last subpath are RW from a fixed interior point until the exit time of a certain domain (between σ and τ we live in the ϵ' interior of Ω, between τ and σ we live in a rectangular band with a flat reflecting boundary). The last subpath is a RW path between two interior points. If we renormalize this to a probability measure, it converges to the corresponding measure for BM. On the other hand, the total mass of this type of subpaths is the expected number of visits of the endpoint from the starting point for a RW stopped upon exiting the domain, i.e. it is the value of the Green's function (with explicit boundary conditions, the normalization is such that the value of the Laplacian at the singularity is 1). Compare this to the continuous setup, where the corresponding mass of Brownian path is given by the time integral of the value of the heat kernel at the origin, which is some measure of the expected time spent by a BM at its starting point before exiting the domain. This time integral of the heat kernel is the Green function. So the rooted ϵ, ϵ'-admissible loops converge. To get the unrooted measure, we divide by the total time and sum or integrate. Note that $1/n \sum_{\Omega^\delta} \cdot = 1/T \sum \delta^2 \cdot \to 1/T \int_\Omega \cdot$. And so the discrete and continuous summations do indeed correspond. □

Acknowledgements

We acknowledge support and/or hospitality of the following grants and institutions: The Einstein Foundation Berlin, TU Berlin, SNF-155922, NCCR Swissmap and the FIM at ETH Zürich, the ESPRC under grant EP/103372X/1, and the Isaac Newton Institute. L.D. also thanks Gábor Pete for useful discussions.

References

[1] Tom Alberts, Michael J. Kozdron, and Robert Masson. Some partial results on the convergence of loop-erased random walk to SLE(2) in the natural parametrization. *J. Stat. Phys.*, 153(1):119–141, 2013.

[2] Martin Barlow. Loop erased walks and uniform spanning trees. MSJ Memoirs, to appear.

[3] Vincent Beffara. The Hausdorff dimensions of SLE curves, *Ann. Probab.* 36, 1421-1452, 2008.

[4] Christian Beneš, Fredrik Johansson Viklund, and Gregory F. Lawler. Scaling limit of the loop-erased random walk Green's function. *Probab. Th. Rel. Fields* 166, 271-319, 2014.

[5] Itai Benjamini. Large scale degrees and the number of spanning clusters for the uniform spanning tree. In *Perplexing problems in probability*, volume 44 of *Progr. Probab.*, 175–183, 1999.

[6] Itai Benjamini, Russell Lyons, Yuval Peres, and Oded Schramm. Uniform spanning forests. *Ann. Probab.*, 29(1):1–65, 2001.

[7] Dmitry Chelkak and Stanislav Smirnov. Discrete complex analysis on isoradial graphs. *Adv. Math.*, 228(3):1590–1630, 2011.

[8] Christophe Garban, Gábor Pete, and Oded Schramm. Pivotal, cluster and interface measures for critical planar percolation, J. Amer. Math. Soc. 26, 939–1024, 2013.

[9] Christophe Garban, Gábor Pete, and Oded Schramm. The scaling limits of near-critical and dynamical percolation, preprint, 2013.

[10] Christophe Garban, Gábor Pete, and Oded Schramm. The scaling limits of the Minimal Spanning Tree and Invasion Percolation in the plane, preprint, 2013.

[11] Geoffrey Grimmett. The random cluster model, Springer, 2006.

[12] Gady Kozma. The scaling limit of loop-erased random walk in three dimensions, *Acta Math.* 191, 29 -152, 2007.

[13] Richard Kenyon. The asymptotic determinant of the discrete Laplacian, *Acta Math.* 185, 239-286, 2000.
[14] Gregory F. Lawler. Intersections of random walks. Birkhäuser, 1991.
[15] Gregory F. Lawler and Vlada Limic. Random walks. A modern introduction, CUP, 2010.
[16] Gregory F. Lawler and Mohammad A. Rezaei. Minkowski content and natural parameterization for the Schramm-Soewner evolution. *Ann. Probab.*, 43:1082-1120, 2015.
[17] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32(1B):939–995, 2004.
[18] Gregory F. Lawler and José A. Trujillo Ferreras. Random Walk Loop Soups. *Trans. Amer. Math. Soc.*, 359(2):767–787, 2007.
[19] Gregory F. Lawler and Fredrik Viklund. Convergence of loop-erased random walk in the natural parametrization, preprint 2016.
[20] Russel Lyons and Yuval Peres. Probability on Trees and Networks, CUP, to appear.
[21] Robert Masson. The growth exponent for planar loop-erased random walk. *Electron. J. Probab.*, 14:no. 36, 1012–1073, 2009.
[22] Oded Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000.
[23] Oded Schramm and Stanislav Smirnov. On the scaling limits of planar percolation, *Ann. Probab.* 39, 1768-1814, 2011.
[24] Stanislav Smirnov. Discrete complex analysis and probability, in *Proceedings of the International Congress of Mathematicians (ICM), Hyderabad, India*, 595-621, 2010.
[25] Wendelin Werner. A simple renormalization flow for FK-percolation models, In the *Jean-Michel Bismut 65th anniversary volume*, to appear.
[26] Wendelin Werner. Topics on the two-dimensional Gaussian Free Field, *Lecture Notes of the topics course at ETH Zurich, September-December 2014*, available online.
[27] David B. Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, 296–303, 1996.
[28] Jean Zinn-Justin, Phase transitions and renormalization group, OUP, 2013.

Department of Mathematics, Columbia University, 2990 Broadway, New York, NY 10027, USA
E-mail address: `sbenoist@math.columbia.edu`

Statistical Laboratory, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK
E-mail address: `L.Dumaz@statslab.cam.ac.uk`

D-Math, ETH Zürich, Rämistr. 101, 8092 Zürich, Switzerland
E-mail address: `wendelin.werner@math.ethz.ch`